Adaptive Quantum-Safe Cryptography for 6G Vehicular Networks via Context-Aware Optimization

Poushali Sengupta
University of Oslo
poushals@ifi.uio.no

Mayank Raikwar
University of Oslo
mayankr@ifi.uio.no

Sabita Maharjan
University of Oslo
sabita@ifi.uio.no

Frank Eliassen
University of Oslo
frank@ifi.uio.no

Yan Zhang
University of Oslo
yanzhang@ifi.uio.no

***Abstract*—Powerful quantum computers in the future may be able to break the security used for communication between vehicles and other devices (Vehicle-to-Everything, or V2X). New security methods called post-quantum cryptography can help protect these systems, but they often require more computing power and can slow down communication, posing a challenge for fast 6G vehicle networks. In this paper, we propose an adaptive post-quantum cryptography (PQC) framework that predicts short-term mobility and channel variations and dynamically selects suitable lattice-, code-, or hash-based PQC configurations using a predictive multi-objective evolutionary algorithm (APMOEA) to meet vehicular latency and security constraints. However, frequent cryptographic reconfiguration in dynamic vehicular environments introduces new attack surfaces during algorithm transitions. A secure monotonic-upgrade protocol prevents downgrade, replay, and desynchronization attacks during transitions. Theoretical results show decision stability under bounded prediction error, latency boundedness under mobility drift, and correctness under small forecast noise. These results demonstrate a practical path toward quantum-safe cryptography in future 6G vehicular networks. Through extensive experiments based on realistic mobility (LuST), weather (ERA5), and NR-V2X channel traces, we show that the proposed framework reduces end-to-end latency by up to 27%, lowers communication overhead by up to 65%, and effectively stabilizes cryptographic switching behavior using reinforcement learning. Moreover, under the evaluated adversarial scenarios, the monotonic-upgrade protocol successfully prevents downgrade, replay, and desynchronization attacks.**

I. Introduction

As large-scale quantum computers become more prevalent, they pose a serious threat to the security of modern vehicle communication systems. Current security measures for Vehicle-to-Everything (V2X) communication mostly rely on techniques like elliptic-curve and RSA-based public-key infrastructures. Unfortunately, these methods are not effective enough, such as those using Shor's algorithm [1]. For critical applications in future 6G vehicle networks, where safety and reliability are essential, switching to quantum-safe security is necessary. Standardized PQC algorithms include lattice-based schemes (e.g., Kyber and Dilithium) [2], [3], code-based schemes (e.g., Classic McEliece) [4], [5], and hash-based signatures (e.g., SPHINCS+) [6]. While PQC algorithm hold great potential, directly deploying PQC in vehicular systems remains challenging. Many PQC algorithms require larger keys, more processing power, and increased computation time resulted in communication delay [7], [8], which can interfere with real-time requirements such as collision avoidance and cooperative vehicle communication. Additionally, traditional models for implementing PQC often overlook the dynamic nature of vehicular contexts [9], [10]. Vehicles face changing conditions, including varying mobility, weather-related impacts on communication quality, and fluctuating computing resources. Depending on these conditions, the demands can range from urgent safety notifications to routine data transfers. A static post-quantum cryptographic configuration cannot simultaneously satisfy all vehicular operating conditions, as overly conservative choices introduce unnecessary delay, while lightweight configurations may offer insufficient security under challenging or adversarial scenarios. Vehicular links exhibit high Doppler spread, frequent beam misalignment, and rapid channel variations [11], [12], leading to transient cost differences between PQC families. Code-based schemes excel during deep fades due to robustness [13], while lattice-based schemes perform better in moderate-SNR, low-latency scenarios [14]. Because these fluctuations occur within sub-100 ms [15], static PQC assignments can cause latency or security overhead, underscoring the need for a predictive, context-aware cryptographic layer that optimizes PQC choices to maintain URLLC compliance [16].

To tackle these issues, we propose a Context-Aware Adaptive PQC (CAAP) Framework tailored for 6G vehicular networks. Unlike conventional, fixed PQC deployments, CAAP dynamically selects the best cryptographic algorithm based on current conditions using real-time sensing and predictive analysis. The framework consists of four main components: (1) a *Context-Sensing Pipeline* that collects vehicle speed, communication quality, weather conditions, and message urgency into a unified context vector; (2) a *Short-Term Predictor* that anticipates context changes in 100–200 ms windows; (3) an *Adaptive Predictive Multi-Objective Evolutionary Algorithm (APMOEA)* that balances latency, compute cost, communication overhead, and quantum-resilience by selecting among lattice-based, code-based, or hash-based signatures; and (4) a *Secure Transition Protocol* that prevents downgrade, replay, and context-manipulation

Workshop on Security and Privacy of Next-Generation Networks (FutureG) 2026
23 February 2026, San Diego, CA, USA
ISBN 979-8-9919276-9-7
https://dx.doi.org/10.14722/futureg.2026.23099
www.ndss-symposium.org

attacks through authenticated version-monotonic negotiation, following principles similar to TLS 1.3 and QUIC [17], [18].

At the core, APMOEA uses reinforcement learning to continually adapt and reduce processing and communication delays while maximizing security. Learning from real-time and past performance, it ensures effective PQC decisions across diverse vehicular environments. We rigorously analyze this adaptive layer, deriving stability conditions for context-aware decisions, bounded latency under mobility drift, and correctness of PQC choice under small prediction error. Through extensive testing under realistic conditions, including LuST mobility traces, ERA5 weather data, 3GPP-compliant NR-V2X channel models, and automotive compute models, our adaptive framework demonstrates a 27% latency reduction, full downgrade-attack resistance, and improved robustness over NSGA-II and RL-only baselines. Main contributions in this paper are as follows:

- **Adaptive planning for PQC selection under vehicular constraints.** We propose APMOEA, a predictive multi-objective planning mechanism that dynamically selects PQC configuration profiles under stringent URLLC latency, computation, communication, and security constraints.
- **Secure execution via monotonic PQC transitions.** We design a lightweight, authenticated transition protocol that enforces monotonic PQC upgrades and prevents downgrade, replay, and desynchronization attacks during reconfiguration.
- **Formal guarantees for stability and bounded latency.** We provide theoretical analysis establishing decision stability under bounded prediction error, correctness relative to an oracle selector, and bounded end-to-end latency under realistic vehicular mobility and channel drift.
- **Claim-driven evaluation on realistic V2X traces.** We conduct experiments using LuST mobility, ERA5 weather data, NR-V2X channel models, and automotive-grade compute profiles, demonstrating latency reductions of up to 27% and robust behavior under adversarial context manipulation.

We emphasize that this work does not introduce new cryptographic primitives. Instead, this work primarily addresses the problem of *adaptive selection and secure transition of post-quantum cryptographic (PQC) mechanisms under stringent vehicular ultra-reliable low-latency communication (URLLC) constraints*. While the system is designed to be compatible with emerging 6G-ready V2X environments, the central contribution lies in *context-aware cryptographic adaptation* rather than detailed physical-layer modeling of future 6G networks. Multi-objective optimization and reinforcement learning are employed as enabling mechanisms to realize this adaptive security objective efficiently and stably. As fully standardized 6G vehicular communication stacks are not yet available, we adopt a *6G-ready* perspective. Specifically, we model vehicular communications that reflect anticipated 6G characteristics, such as URLLC, high mobility, and dense connectivity, while remaining compatible with current NR-V2X abstractions. **Scope of PQC deployment:** CAAP applies post-quantum cryptography at the application layer to protect V2V data communications transmitted over established 5G connections. The PQC schemes encrypt and authenticate safety-critical and cooperative awareness messages exchanged between vehicles. This is distinct from and complementary to the network-layer security provided by 5G-AKA and EAP-TLS, which handle initial authentication and secure connection establishment between vehicles and the 5G network infrastructure. [1]

II. Background, Related Work, and Threat Model

This section covers the background, prior work, and adversarial assumptions for adaptive PQC in V2X systems.

A. Background and Related Work

Next-generation vehicular network services, which will be developed on beyond 5G and 6G technologies demand ultra-reliable low-latency communication (URLLC) with latency requirements between 5 and 20 milliseconds, high reliability, and the ability to operate effectively under rapidly changing mobility and channel conditions [19], [20]. Traditional security methods based on Elliptic Curve Cryptography (ECC) and RSA are vulnerable to quantum computing threats, making the adoption of PQC crucial for future implementations [21], [22]. However, when deployed on today's classical vehicular hardware, post-quantum cryptographic schemes generally incur larger key sizes, higher verification costs, and increased computational overhead compared to traditional public-key cryptography, all of which can fluctuate significantly with factors such as SNR, packet error rate (PER), CPU load, and changes in mobility [23]–[25]. As a result, static PQC configurations often violate URLLC latency constraints under strict operating conditions or introduce avoidable computational and communication overhead during low-risk scenarios.

Most recent proposals for PQC in V2X systems mainly evaluate algorithms such as Kyber, Dilithium, and McEliece under fixed conditions [26], [27]. These evaluations do not account for the effects of rapid vehicle movement, the fading processes described by 3GPP, or variations in available computational resources. Additional benchmarking studies [28] has further demonstrated that no single PQC method outperforms others across all operational scenarios. This highlights the necessity for a PQC layer that can adapt based on predictions and contextual awareness. While adaptive and resource-aware cryptographic strategies for the Internet of Things (IoT) and embedded systems [29]–[31] adjust classical key strengths according to load or energy availability, these approaches do not handle significant variability in PQC parameters nor do they protect against downgrade, rollback, or context-forcing attacks that may occur during PQC transitions. Protocols like TLS 1.3 and QUIC [32], [33] provide mechanisms to resist downgrade attacks, but assume stable client-server environments, unlike the dynamic conditions in V2X scenarios, which involve unpredictable latency. Prior research on detecting downgrades [34] and negotiating QUIC versions [35] has highlighted the challenges of preventing rollback in hostile network environments, but has not integrated real-time context aware adaptation. To the

[1]This work is supported by the RCN Transport 2025 project CRISP (No. 302327).

best of our knowledge, CAAP is the first to merge downgrade-resistant PQC transitions with a reinforcement-learning-guided predictive optimizer, specifically tailored for the demands of 6G V2X operations. This work does not introduce new post-quantum cryptographic primitives; instead, it focuses on the adaptive orchestration of standardized PQC mechanisms (e.g., Kyber, Dilithium, SPHINCS+) in response to rapidly changing vehicular contexts.

B. Threat Model

We consider a powerful adversary with access to V2X wireless communications who can inject, replay, delay, or modify packets [32], [36]. The attacker may forge contextual features such as SNR or PER to influence PQC selection [37], replay outdated version counters to trigger downgrade attacks [33], or manipulate message delivery to induce synchronization failures [37]. The adversary's objectives include weakening PQC strength, increasing latency, destabilizing transitions, or steering the system toward slower or less secure configurations. In practice, post-quantum security relies on coordinated cryptographic mechanisms rather than a single primitive. Each candidate action therefore represents a *PQC configuration profile* comprising (i) a key encapsulation mechanism (KEM) for session key establishment and (ii) a digital signature scheme for authentication and integrity. For example, configurations may combine Kyber for key exchange with Dilithium or SPHINCS+ for authentication. We assume standardized, NIST-recommended PQC primitives and do not modify underlying cryptographic algorithms. In addition, the adversary may degrade channel quality [38]. Vehicles and roadside units (RSUs). RSUs are not mandatory in 5G NR-V2X or future 6G architectures. Our framework fully supports direct V2V communication, treating RSUs as optional edge infrastructure that can reduce latency when available, are assumed to possess trusted hardware for secure long-term key storage and reliable version counters [39], and PQC implementations are assumed free of side-channel leakage. While onboard sensor data (e.g., speed and environmental signals) are trusted, network-observable metrics may be manipulated. CAAP enforces (i) confidentiality, (ii) message integrity and authenticity, (iii) monotonic, non-reversible PQC upgrades, and (iv) consistent PQC transitions across communicating endpoints.

We address a strong adversary model in vehicular security, featuring a quantum-capable attacker, an active network adversary, a context-manipulation attacker, and a downgrade attacker. While CAAP relies on trusted vehicle hardware, it operates over an untrusted communication medium, ensuring quantum-resistant confidentiality, integrity, and reliable upgrades despite these threats. The main goal of this work is to enhance PQC schemes by reducing delays. However, understanding security threats is vital for two reasons. First, adaptable cryptographic systems can be targeted by attackers who may manipulate their environment to force the use of weaker algorithms. Second, switching between cryptographic methods is a critical process; if an attacker disrupts this, they can compromise both security and performance. Thus, it's crucial to evaluate the optimization engine and secure transition protocols against potential attacks to ensure the reliability of the low-latency PQC framework in 6G vehicular environments.

TABLE I: System Assumptions and Adversary Capabilities.

Category	Description
Assumptions	• Trusted hardware roots-of-trust (secure key storage, monotonic counters), authenticated boot, and correct PQC implementations. Vehicles and RSUs are uncompromised. • The wireless channel is untrusted and effectively adversary-controlled. • Sensors generating context (speed, SNR, mobility) are noisy but not maliciously compromised.
Adversary Capabilities	• Quantum-capable attacker able to break classical cryptosystems (ECDSA/ECDH/RSA). Harvest-now–decrypt-later is allowed. • Active network attacker who can inject, modify, replay, jam, or drop V2X messages. • Context-manipulation attacker who induces artificial channel degradation or packet loss to influence PQC selection. • Downgrade/rollback attacker manipulating transition messages to enforce weaker PQC.

III. Context aware Adaptive PQC (CAAP) Framework

We present CAAP that adjusts the type of PQC scheme used based on the current situation of vehicles. The system consists of four cooperating components: (i) **context extraction** for mobility, channel, and workload indicators; (ii) **short-term prediction** of context evolution using lightweight forecasting; (iii) **APMOEA**, which selects the optimal PQC algorithm; and (iv) a **secure transition protocol** that enforces authenticated, monotonic version upgrades. The context is updated every 20–50 ms, predictions operate over 100–200 ms horizons, and PQC transitions complete within a single V2X round-trip. Vehicles actively monitor a range of factors that influence the cost and speed of cryptographic processes. These include mobility features (such as speed and connection duration), channel conditions (such as SNR and PER), environmental factors (including weather and visibility), computational load (referring to CPU and GPU resources), and message urgency (differentiating between critical safety communications and less urgent messages). These various inputs are standardized and sent to the next module for processing. To prevent rapid changes from causing erratic shifts in cryptographic methods, we utilize short-term predictive modeling. This involves using simple filters and regression models to predict what will happen next (such as expected latency, SNR, and processing capabilities). This forecasting enables us to provide stable data for decision-making while keeping computational demands low. The decision engine determines the best PQC algorithm to use during each communication session. The potential algorithms include: **1. Lattice-based schemes** (like Kyber and Dilithium), which balance speed and robust, **2. Code-based schemes** (such as Classic McEliece), ideal for use in poor communication conditions, and **3. Hash-based signatures** (e.g., SPHINCS+), which are lightweight and work without maintaining state. Every algorithm is evaluated based on a multi-dimensional cost vector, allowing for an informed selection to meet the specific needs of the situation, which is:

$$C_{\text{alg}} = (T_{\text{enc}},\ T_{\text{dec}},\ S_{\text{key}},\ S_{\text{ct}},\ E_{\text{comp}},\ S_{\text{sig}}), \tag{1}$$

representing encryption and decryption time, key sizes, ciphertext size, computational energy, and signature size. These vectors are input for the optimization described in Section IV.

Note on evaluation scope: In this work, we evaluate encryption/key encapsulation schemes (Kyber, Classic McEliece) and digital signature schemes (Dilithium, SPHINCS+) individually to assess their operational feasibility and robustness under V2V network constraints. We acknowledge that real-world deployments require both mechanisms working in tandem—encryption to defend against eavesdropping and signatures to prevent active man-in-the-middle attacks. While our adaptive selection currently operates on individual primitives, the framework is designed to support hybrid configurations.

A. Relation to the MAPE-K Adaptation Loop

The proposed PQC upgrade mechanism does not modify existing cryptographic standards; it operates at the protocol orchestration level using standardized, NIST-recommended PQC algorithms, following downgrade-protection principles similar to TLS 1.3 and QUIC. The impact on application throughput is captured indirectly through communication overhead and end-to-end latency under URLLC constraints. While evaluated in a V2X setting, the framework is not inherently domain-specific and can be extended to other latency-sensitive systems.

CAAP follows the classical *Monitor-Analyze-Plan-Execute over Knowledge (MAPE-K)* control loop, widely used in self-adaptive and autonomic systems [40]. Figure 1 illustrates the CAAP organized according to the MAPE-K control loop, used as an organizing abstraction to improve interpretability. The distributed **knowledge base (K)** comprises offline PQC cost profiles, reinforcement-learning feedback, and protocol state (e.g., version counters and context hashes) supporting robust adaptation. This work does not propose a new adaptation architecture; instead, it contributes a predictive multi-objective planning formulation for PQC selection with stability guarantees, together with a secure, monotonic execution protocol that prevents downgrade and desynchronization—capabilities absent from generic self-adaptive frameworks.

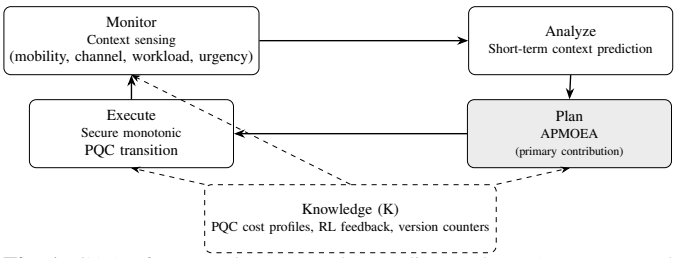

Fig. 1: CAAP framework structured according to the MAPE-K control loop.

B. V2X System Model and Vehicular Dynamics

The vector C_{alg} captures the intrinsic, context-independent cost characteristics of a given PQC configuration profile, obtained through offline benchmarking. At runtime, these algorithm-level properties are mapped to context-dependent objectives that reflect vehicular communication conditions, system load, and message urgency. Accordingly, the optimization engine operates on a derived multi-objective function rather than directly on C_{alg}. We consider a vehicular communication environment based on NR-V2X abstractions, operating under ultra-reliable low-latency communication (URLLC) constraints. Vehicles exchange periodic and event-driven messages with roadside units (RSUs) and neighboring vehicles in the presence of high mobility and rapidly varying channel conditions. Vehicular mobility is characterized by time-varying speed $v(t)$ and acceleration $a(t)$, which jointly affect Doppler spread and channel coherence time. To capture short-term link stability without explicit physical-layer modeling, we introduce a *connectivity horizon* τ_c, representing the expected duration for which the communication link remains reliable before a significant topology or channel change occurs. Let $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$ denote the set of candidate cryptographic actions available to the system. Each action $a \in \mathcal{A}$ corresponds to a valid PQC configuration profile rather than an individual cryptographic primitive. Messages are categorized into urgency classes $\mathcal{M}$={safety, control, telemetry}. Safety messages are delay-critical and must satisfy strict latency bounds, whereas control and telemetry messages tolerate higher cryptographic overhead. Each message class is associated with an urgency weight $u \in \{u_s, u_c, u_t\}$ that influences PQC selection decisions.

C. Context Model and Notation

The transition protocol relies on real-time information from the vehicle's surroundings, the wireless connection, and the device's processing state. We represent this information using a context vector defined as:

$$X_t = (\text{SNR}, \text{PER}, v_s, v_a, \tau_c, u, \text{CPU}_{\text{load}}, \ldots). \tag{2}$$

The connectivity horizon τ_c and urgency weight u allow the optimizer to account for vehicular mobility and safety-critical message semantics without requiring detailed PHY-layer simulation. This context vector includes key factors like signal strength (SNR), packet error rate (PER), the vehicle's speed (v_s), its acceleration (v_a), visibility conditions, ambient temperature (T_{amb}), current CPU load, and how urgent the message is. These elements give us a comprehensive view of the vehicle's physical movement, the quality of the wireless connection, environmental conditions, processing capacity, and the importance of the messages being transmitted. Additionally, all cost functions in the APMOEA optimizer are Lipschitz continuous with respect to these features. This means that even if there are small changes in the context, the system will behave consistently and predictably. Each PQC configuration profile incurs different computational, communication, and latency overheads, which are particularly critical in vehicular URLLC scenarios. For instance, hash-based signature schemes offer strong security guarantees at the cost of increased signature size, whereas lattice-based schemes provide lower latency but rely on different computational assumptions. The proposed adaptive framework selects among these profiles based on real-time vehicular context, message urgency, and resource availability.

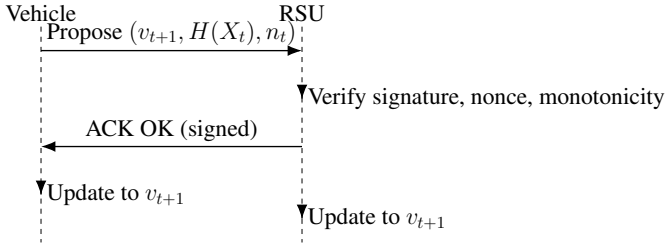

Fig. 2: PQC version-upgrade protocol with monotonic transitions.

D. Secure PQC Transition Protocol

When switching between different PQC configurations, additional security risks can arise during the transition process. To mitigate these risks, we implement a secure negotiation protocol with the following key features: authenticated signaling messages to verify message origin, monotonic versioning to prevent rollback to less secure configurations, atomic transitions to ensure that updates occur consistently, replay protection through unique identifiers, and decision consistency mechanisms to ensure that both communicating parties agree on the selected PQC configuration. Together, these measures protect against transition-phase attacks and preserve secure communication. The secure PQC transition protocol, as illustrated in Figure 2, orchestrates smooth version upgrades and secure negotiation between vehicles and RSUs. The process involves gathering real-time data, predicting short-term conditions, running an optimization algorithm to find the best encryption options, selecting the optimal PQC algorithm, securely negotiating new parameters, and continuously monitoring the situation. This adaptive framework delivers quantum-resistant security while remaining responsive to latency requirements and resource constraints; essential characteristics for 6G-V2X networks.

IV. APMOEA Optimization Algorithm

APMOEA serves as the core decision engine for adaptive PQC configuration selection under latency, computational, and security constraints in dynamic vehicular contexts. To improve decision stability and accelerate convergence under rapidly changing conditions, reinforcement-learning (RL) feedback is integrated as an internal mechanism within the optimization loop. For each candidate PQC configuration profile $a \in \mathcal{A}$, we define a multi-objective cost vector:

$$f(a) = (T_{\text{lat}}(a),\ C_{\text{comp}}(a),\ S_{\text{comm}}(a),\ \sigma_{\text{sec}}(a))\,, \tag{3}$$

where, $T_{\text{lat}}(a)$ is end-to-end signing/verification latency, $C_{\text{comp}}(a)$ is computational cost (CPU cycles), $S_{\text{comm}}(a)$ captures key and payload size overhead, and $\sigma_{\text{sec}}(a)$ is cryptographic strength (bit security). Formally, the objective vector $f(a)$ is computed as a context-aware transformation of C_{alg} given the current state X_t. APMOEA seeks the Pareto-optimal solution set under real-time constraints. Predicted context variables (e.g., channel quality, processing availability) from Section III are incorporated as dynamic weights. The optimizer follows a series of core steps: First, it initializes a population of PQC configurations based on context features. Then, it uses mutation and crossover techniques to explore the design space and generate candidate configurations. These candidates are ranked using a traditional context-weighted functional form of a fitness (utility) function, $\text{Fitness}(a) = w_1 T_{\text{lat}} + w_2 C_{\text{comp}} + w_3 S_{\text{comm}} - w_4 \sigma_{\text{sec}}$, which combines factors like latency, computation costs, communication costs, and security. Finally, a reinforcement learning agent monitors the performance of the selected algorithms, adjusting mutation rates and weight parameters to enhance stability and minimize fluctuations in PQC algorithm choices. Consider an intelligent transportation system managing V2X communications to enhance road safety and efficiency. In this scenario, a fleet of autonomous vehicles is constantly communicating with each other and with traffic infrastructure to make real-time decisions.

1) **State Vector Initialization:** Each vehicle gathers data that forms the state vector s_t. For instance, a vehicle's state might include $T_{\text{lat}}, C_{\text{comp}}, \text{SNR}, \text{mobility}, \text{CPUload}$
2) **Decision-Making and Reward System**: As the vehicles operate, they consistently assess their contexts. For example, High latency (T_{lat}) negatively impacts operations, resulting in a lower reward score R_t, especially if communication channels are switched frequently (N_{switch}). Conversely, a stable connection with low latency and strong SNR leads to higher rewards. Additionally, the reward is enhanced for vehicles with strong cryptographic security (σ_{sec}).
3) **Adapting Strategy**: The vehicle adapts its communication behavior based on observed rewards; when latency increases, it may adjust transmission parameters or select alternative PQC configuration profiles to improve performance.
4) **Continuous Learning Process**: Vehicles continuously adapt to various road conditions, traffic situations, and environmental changes, refining their strategies based on feedback to remain efficient in heavy traffic or adverse weather.

We employ a reward function that minimizes latency, penalizes excessive switching and poor SNR conditions, and incentivizes the selection of higher-security PQC configurations as follows:

$$R_t = -T_{\text{lat}} - \alpha \cdot N_{\text{switch}} + \beta \cdot \text{SNR} + \gamma \cdot \sigma_{\text{sec}}\,. \tag{4}$$

We employ a tabular Q-learning approach to avoid the computational overhead associated with neural network–based reinforcement learning. APMOEA converges to stable solutions by adapting dynamic weights to predicted conditions, penalizing oscillatory decisions with reinforcement learning feedback, and gradually reducing exploration as stable fitness patterns emerge. These features enable the algorithm to consistently generate PQC recommendations in the fast-changing conditions of vehicular environments. The adaptive optimization process is formalized in Algorithm 1. APMOEA incorporates short-term context predictions by dynamically adjusting the relative weights of latency, computation, communication, and security objectives prior to fitness evaluation. In each generation, the algorithm evaluates a fixed-size population of candidate PQC configuration profiles, ensuring bounded and predictable computational cost. For four PQC families and population sizes below 40, the optimization completes in under 1 ms

on automotive-grade CPUs, enabling real-time cryptographic selection in vehicular environments.

Algorithm 1 Adaptive Predictive Multi-Objective Evolutionary Algorithm (APMOEA)

Input: Current context X_t, predicted context $\hat{X}_{t+1}$, PQC algorithm set $\mathcal{A}$
Output: Selected PQC algorithm a_t for the current context
Initialize population P_0 by sampling candidate algorithms from $\mathcal{A}$
for $t = 1$ to T **do**
  Compute multi-objective cost vector $\mathbf{f}(a)$ for each $a \in P_t$
  Update dynamic weights w_i based on predicted context $\hat{X}_{t+1}$
  Compute fitness scores for all candidates using the weighted cost
  Select parent candidates using tournament selection (pairwise comparison of randomly sampled candidates)
  Generate offspring through crossover and mutation
  RL agent adjusts mutation rate and weight parameters to reduce instability
  Form the next population P_{t+1} from parents and offspring
end for
Identify the Pareto-optimal candidate (non-dominated by any other) and output as a_t

Note: *Tournament selection means that a small random subset of candidates is sampled, and the one with the highest fitness in that subset is chosen as a parent. A Pareto-optimal solution is a candidate for which no other algorithm performs better in all objectives simultaneously.*

In APMOEA, the population consists of candidate PQC configuration profiles. Each profile specifies a standardized PQC primitive (e.g., Kyber, Dilithium, McEliece, or SPHINCS+) together with a weight vector that reflects the relative importance of latency, computation cost, communication overhead, and security. During optimization, evolutionary operators are applied to these profiles. Crossover combines the weight vectors of well-performing configurations to create new trade-offs, while mutation introduces small random changes to further explore nearby preferences. These operations do not generate new cryptographic schemes; instead, they produce new preference profiles that guide the selection among existing PQC primitives. This enables the optimizer to adaptively choose the most suitable PQC configuration as vehicular conditions change, while preserving the correctness and integrity of the underlying cryptographic algorithms.

Algorithm 2 Secure PQC Transition Protocol

Input: Current version v_t, proposed version v_{t+1}, context hash $H(X_t)$
Output: Parties update to version v_{t+1}
Sender constructs transition message:

$$M = \big(v_{t+1}, H(X_t), \text{nonce}\big)$$

Sender signs M using current scheme v_t
Receiver verifies signature and checks:
- $v_{t+1} \geq v_t$ (monotonic check)
- $H(X_t)$ matches its local context hash
- nonce has not been used before

if any check fails **then**
  Reject transition and revert to v_t
else
  Receiver acknowledges with signed confirmation
  Both parties update to version v_{t+1}
end if

The secure PQC transition protocol allows vehicles and roadside units (RSUs) to safely upgrade to new post-quantum cryptographic configurations without risking downgrade, replay, or desynchronization attacks. At each update, the sender proposes a new PQC version together with a context hash and a fresh nonce, and signs this message using the currently active PQC scheme. The receiver verifies that the proposed version is not lower than the current one, that the context hash matches its local view, and that the nonce is fresh. If any check fails, the transition is rejected and both parties continue using the existing configuration; otherwise, a signed acknowledgment is exchanged and both sides switch to the new version simultaneously.

V. Theoretical Guarantees

In this section we analyze the stability, robustness, and security properties of the CAAP. We show that APMOEA consistently selects appropriate PQC configuration profiles under dynamic vehicular conditions, while the secure transition protocol strengthens cryptographic guarantees during reconfiguration. Observed increases in end-to-end latency remain bounded and within URLLC budgets relevant to vehicular dynamics, including channel coherence time and mobility-induced variability. Furthermore, the optimizer maintains robust selection behavior as long as short-term mobility and channel prediction errors remain within moderate bounds. Let X_t denote the contextual feature vector at time t, and let $\hat{X}_{t+1}$ be its predicted value. APMOEA selects a PQC algorithm through, $a_t = \arg\min_{a \in \mathcal{A}} L_t(a)$, where the weighted cost is, $L_t(a) =$

$$w_1\, T_{\text{lat}}(a, X_t) + w_2\, C_{\text{comp}}(a, X_t) + w_3\, S_{\text{comm}}(a) - w_4\, \sigma_{\text{sec}}(a). \quad (5)$$

We assume: **(A1)** all cost terms are L-Lipschitz in context X_t; **(A2)** mobility and channel drift satisfy $\|X_{t+1} - X_t\| \leq \delta$; **(A3)** prediction errors satisfy $\|\hat{X}_{t+1} - X_{t+1}\| \leq \varepsilon$; **(A4)** the minimum loss gap between PQC candidates is $\Delta_{\min} > 0$. Let $\Delta_{\min}$ be the minimum loss separation between any two PQC algorithms. **Proofs for the theorems are in the Appendix.**

Theorem V.1 (Decision Stability). *If $K\varepsilon < \Delta_{\min}$, then APMOEA selects the same algorithm a_t for all t within a context-stable interval. Here, K denotes the Lipschitz constant of the loss function, ε is the context-prediction error, and $\Delta_{\min}$ is the minimum loss gap between any two PQC algorithms.*

The protocol prevents unstable PQC switching caused by sensor noise or small prediction errors, helping maintain stable vehicle operation. When prediction errors are sufficiently small, the optimizer makes consistent decisions and the selected PQC configuration remains stable over time. Formally, stability holds when $K\,\varepsilon < \Delta_{\min}$. We evaluate robustness against rollback attempts, message loss, version-counter tampering, and synchronization errors, with Table II summarizing defenses.

TABLE II: Failure-mode analysis of PQC transition protocol.

Attack Scenario	Mitigation Mechanism	Outcome
Replay of old version	Nonce + context hash	Detected
Message loss	Atomic 2-phase confirmation	No partial update
Asymmetric update	Mutual acknowledgment	Safe rollback to v_t
Counter modification	Signature mismatch	Packet dropped
Forced downgrade	Monotonic version check	Impossible

The mitigation mechanisms in Table II are implemented at the protocol level and are independent of the specific PQC configuration selected by APMOEA. The underlying post-quantum cryptographic primitives are used without modification. Security during reconfiguration is ensured through protocol mechanisms such as authenticated version counters, contextual hashes, freshness nonces, and atomic confirmation steps, which prevent

replay, downgrade, and desynchronization attacks. APMOEA is responsible for selecting PQC configuration profiles, while the transition protocol securely manages the switch between them. A downgrade attempt is detected through version mismatch, resulting in successful acceptance for legitimate upgrades and rejection for downgrade attacks. Let v_t denote the active PQC version at time t and v_{t+1} the proposed version; all transition messages include authenticated version information, context hashes, and freshness nonces to enforce secure transitions.

Theorem V.2 (**Monotonic Upgrade Security**). *Let v_t denote the current PQC version and v_{t+1} the negotiated target version. Under authenticated signaling with version counters, contextual hashes, and mutual confirmation, no probabilistic polynomial-time adversary can induce a transition to any $v < v_{t+1}$ or cause endpoint desynchronization.*

Here, v_t is the PQC version currently in use and v_{t+1} is the new version proposed during the upgrade. Downgrades fail because any tampering breaks the signed, nonce-bound transition record, causing both parties to immediately reject it. This eliminates the most serious risk during hybrid PQC migration. The end-to-end PQC latency is modeled as,

$$T_{\text{lat}}(a, X_t) = T_{\text{enc}}(a) + T_{\text{dec}}(a) + T_{\text{net}}(a, X_t), \tag{6}$$

where T_{net} captures SNR, bandwidth, & load-dependent effects.

Theorem V.3 (**Latency Boundedness**). *Let X_t denote the context at time t. Assume bounded vehicular and channel variations over one decision interval: $|v_{t+1} - v_t| \le \Delta v, \quad |\gamma_{t+1} - \gamma_t| \le \Delta\gamma$, where v_t is vehicle speed and γ_t denotes the link SNR (or a channel-quality proxy). The latency of the PQC configuration selected by APMOEA satisfies*

$$T_{\text{lat}}(a_t, X_t) \le \max_{a \in \mathcal{A}} T_{\text{lat}}(a, X_t) + O(\Delta v + \Delta\gamma). \tag{7}$$

Here, X_t denotes the context at time t, δ bounds the per-step context variation, $a_t \in \mathcal{A}$ is the PQC algorithm selected by APMOEA, and $T_{\text{lat}}(\cdot)$ denotes end-to-end latency. Latency remains bounded since small context changes induce limited latency variation, and APMOEA never selects a candidate exceeding the feasible maximum. Thus, PQC-induced latency remains URLLC-compliant even under rapid mobility. Define the oracle-optimal algorithm under the true next context as $a^\star = \arg\min_{a \in \mathcal{A}} L_{t+1}(a)$, corresponding to the ideal choice with perfect knowledge of the next context.

VI. Experimental Evaluation

We evaluate CAAP against the claims made in our contributions. We evaluate (C1) latency reduction by comparing APMOEA against static PQC baselines under realistic NR-V2X mobility, channel, weather, and compute traces; (C2) downgrade and desynchronization resistance by subjecting the secure transition protocol to replay, rollback, and context-manipulation attacks; (C3) decision stability by introducing controlled prediction noise and measuring PQC switching behavior; and (C4) robustness by analyzing sensitivity to heterogeneous channel conditions and automotive compute profiles.

1. Experimental Setup All experiments are designed to directly evaluate one or more of the stated claims. We consider four NIST-standardized PQC schemes: Kyber-768 and Dilithium-3 (lattice-based), Classic McEliece-348864 (code-based), and SPHINCS+-128s (hash-based). Their computational and communication characteristics are summarized in Table III. APMOEA is compared against three static baselines commonly used in V2X systems: Static-Lattice, Static-Code, and Static-Hash. Evaluation focuses on four metrics: end-to-end latency, computational overhead, communication overhead, and security consistency (i.e., absence of downgrade or desynchronization). Context signals are time-aligned; prediction, optimization, and PQC switching execute on an automotive-grade 800,MHz CPU simulator. Identical context evolution ensures fair comparison.

2. Results & Discussion Experiments are conducted using LuST mobility traces, ERA5 weather data, and NR-V2X channel models, with cryptographic costs derived from NIST PQC reference implementations and automotive-grade CPU profiles.

Latency and URLLC compliance: Adaptive PQC selection consistently reduces end-to-end cryptographic latency compared to static PQC configurations (Table IV), lowering CPU usage by approximately 22–40% through avoidance of computationally intensive schemes under adverse conditions. While classical cryptography (ECDH/ECDSA) exhibits lower latency due to smaller keys and lighter verification, fixed PQC configurations frequently violate the 5–20,ms URLLC budget for V2X safety traffic. In contrast, CAAP maintains latency within URLLC limits across dynamic contexts, demonstrating practical feasibility of post-quantum security for vehicular systems.

Robustness under context manipulation: Figure 3(a) shows that adversarial manipulation of context signals increases the absolute cost of PQC configurations but preserves their relative ordering. This indicates that bounded perturbations do not induce erratic or adversary-controlled cryptographic selection, supporting the robustness assumptions underlying the cost model and optimization process.

Selection behavior and decision stability: Figure 3(b) compares selection behavior across NSGA-II, RL-only, and APMOEA-based strategies. NSGA-II exhibits limited adaptation and collapses to a dominant configuration, while RL-only approaches overreact under low-SNR conditions, frequently selecting high-overhead McEliece-based schemes. In contrast, APMOEA achieves balanced selection across PQC families by jointly optimizing latency, computation, communication, and security. Reinforcement learning further stabilizes decisions, reducing PQC switching frequency from 14.2 to 4.3 switches per minute and lowering peak latency by over 75% (Table IV).

Hardware feasibility: Figure 3(c) evaluates runtime sensitivity on automotive-grade hardware. Although McEliece has low computation cost, its communication overhead violates URLLC constraints. SPHINCS+ is more sensitive to CPU frequency due to verification overhead. The adaptive planner captures these trade-offs and selects configurations that meet millisecond-

TABLE III: Comparison of computational cost, communication overhead, dynamic selection frequency, and adversarial sensitivity for PQC configuration profiles.

Algorithm	Compute Cost (ms)			Comm. Overhead (KB)		Dynamic Selection		Context Attack Cost	
	Enc	Dec	Verify	PK Size	Sig/CT	Sel.%	Trigger	Honest	Attack
SPHINCS+-128s	0.90	–	1.10	0.05	17.00	69.1%	Stateless / fallback under instability	1.89	2.15
Kyber-768	0.25	0.35	0.15	1.18	1.08	26.9%	Low-latency and compute-efficient operation	1.09	1.35
Dilithium-3	0.40	0.45	0.30	1.50	2.70	4.0%	High-assurance authentication phases	1.29	1.55
McEliece-348864	0.05	0.06	–	240.00	0.13	<1%	Rare use under extreme noise	1.59	1.85

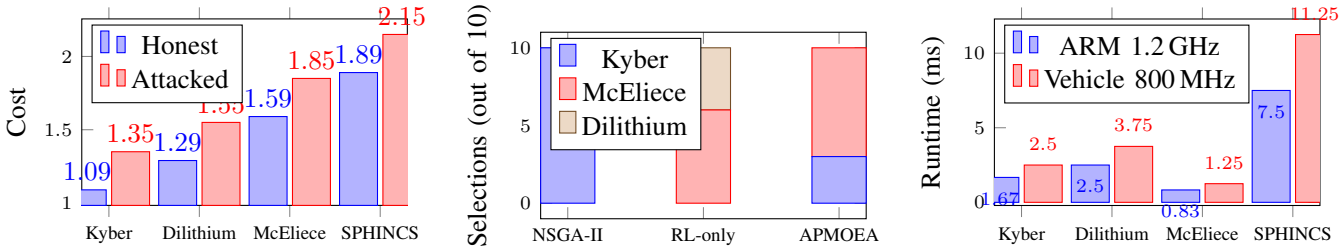

(a) Context-manipulation effect (b) NSGA-II vs RL vs APMOEA (c) Runtime sensitivity across hardware

Fig. 3: (a) PQC cost under adversarial context manipulation, (b) selection behavior across NSGA-II, RL-only, and APMOEA, and (c) hardware-dependent runtime variability (measured in milliseconds) across ARM 1.2 GHz and vehicle-grade 800 MHz CPUs..

TABLE IV: Latency performance, switching stability, and downgrade-attack robustness under realistic NR-V2X conditions.

Scheme / Setting	Key Size (KB)	End-to-End Latency (ms)	Switches / 60s	Security Outcome
ECDH / ECDSA (Classical)	0.09	4.1	–	Classical baseline
Kyber-768	1.18	9.3	–	PQC (static)
Dilithium-3	1.50	10.8	–	PQC (static)
McEliece-348864	240.00	8.7	–	PQC (static)
SPHINCS+-128s	0.05	17.4	–	PQC (static)
Adaptive PQC (APMOEA, no RL)	–	7.8	14.2	Upgrade accepted; downgrade detected
Adaptive PQC (APMOEA + RL)	–	**7.2**	**4.3**	**Upgrade accepted; downgrade detected**

level latency bounds on 800 MHz vehicular CPUs, enabling real-time deployment.

TABLE V: Context manipulation and hardware runtime.

Algorithm	Honest Cost	Attacked Cost	ARM 1.2GHz (s)	Veh. 800MHz (s)
Kyber	1.09	1.35	0.00167	0.00250
Dilithium	1.29	1.55	0.00250	0.00375
McEliece	1.59	1.85	0.00083	0.00125
SPHINCS+	1.89	2.15	0.00750	0.01125

Security of cryptographic transitions: The secure transition protocol successfully prevents downgrade and replay attacks while allowing legitimate PQC upgrades (Table IV, Figure 4). Even under increased packet loss and context noise, reinforcement learning stabilizes planning decisions and prevents harmful or inconsistent transitions. Moreover, PQC selection remains stable for prediction errors up to $\varepsilon \leq 0.10$, consistent with the theoretical stability guarantees derived in Section V.

3. Limitations and Future Scope Our experiments use trace-driven NR-V2X channel models and automotive compute profiles, as 6G systems are unavailable. CAAP supports current 5G NR-V2X and transitions to 6G with minimal changes. Limitations include reliance on clock synchronization, sensitivity to prediction accuracy, and the absence of live deployment validation. Future work will address these via 6G sensing integration, formal verification of RL-assisted adaptation, and end-to-end deployment on solution-based testbeds (e.g., hardware-in-the-loop or small-scale V2X setups) to validate runtime behavior, integration overhead, and system-level interactions beyond trace-based modeling. Future work will extend CAAP to jointly optimize encryption and signature schemes, better reflecting practical post-quantum security where neither mechanism alone is sufficient.

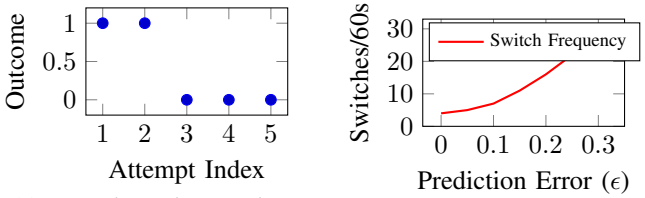

(a) Upgrade vs downgrade (b) Prediction error effect

Fig. 4: (a) downgrade-attack rejection (1=Success, 0=Rejected) and (b) prediction-error-driven switching frequency.

VII. Conclusion

We proposed an adaptive PQC framework CAAP for 6G V2X that jointly predicts mobility and channel dynamics, selects optimal PQC schemes via APMOEA, and enforces secure monotonic upgrades. The framework reduces latency by 27%, prevents downgrade attacks, and remains robust to context manipulation, hardware variability, and prediction noise. Theoretical analysis establishes decision stability and bounded latency under realistic drift, demonstrating a practical path toward quantum-safe vehicular networks.

REFERENCES

[1] T. N. Turnip, B. Andersen, and C. Vargas-Rosales, "Towards 6g authentication and key agreement protocol: A survey on hybrid post quantum cryptography," *IEEE Communications Surveys & Tutorials*, 2025.

[2] L. Chen, D. Moody, and A. Patrick, "Nist pqc round 3 finalists: Performance evaluation," National Institute of Standards and Technology, Tech. Rep., 2021.

[3] M. A. Khan, S. Javaid, S. A. H. Mohsan, M. Tanveer, and I. Ullah, "Future-proofing security for uavs with post-quantum cryptography: A review," *IEEE Open Journal of the Communications Society*, 2024.

[4] R. McEliece, "A public-key cryptosystem based on algebraic coding theory," *DSN Progress Report*, 1978.

[5] N. Sendrier, "Code-based cryptography: State of the art and perspectives," in *IEEE ISIT*, 2011.

[6] D. Bernstein *et al.*, "Sphincs+: Submission to nist pqc standardization," 2019.

[7] A. Nsour, "Quantum-resilient secure onboard communication (qrsecoc): Integrating post-quantum cryptography for robust automotive network security," *INTERNATIONAL JOURNAL*, vol. 12, no. 1, pp. 27–48, 2025.

[8] S. Paul, P. Scheible, and F. Wiemer, "Towards post-quantum security for cyber-physical systems: Integrating pqc into industrial m2m communication," *Journal of Computer Security*, vol. 30, no. 4, pp. 623–653, 2022.

[9] X. Zhang, J. Li, J. Zhou, S. Zhang, J. Wang, Y. Yuan, J. Liu, and J. Li, "Vehicle-to-everything communication in intelligent connected vehicles: A survey and taxonomy," *Automotive Innovation*, pp. 1–33, 2025.

[10] C. Jung, D. Lee, S. Lee, and D. H. Shim, "V2x-communication-aided autonomous driving: System design and experimental validation," *Sensors*, vol. 20, no. 10, p. 2903, 2020.

[11] M. Deng, M. Ahmed, A. Wahid, A. A. Soofi, W. U. Khan, F. Xu, M. Asif, and Z. Han, "Reconfigurable intelligent surfaces enabled vehicular communications: A comprehensive survey of recent advances and future challenges," *IEEE Transactions on Intelligent Vehicles*, 2024.

[12] J. Tan, T. H. Luan, W. Guan, Y. Wang, H. Peng, Y. Zhang, D. Zhao, and N. Lu, "Beam alignment in mmwave v2x communications: A survey," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 3, pp. 1676–1709, 2024.

[13] V. Nguyen, "Code-based cryptography: Attacking and constructing cryptographic systems," Ph.D. dissertation, Lund University, 2025.

[14] W. Ding, "Adaptive methods toward hyper-reliable and low-latency communication," Ph.D. dissertation, King's College London, 2024.

[15] N. Agarwal, "Alternative shapes of modulation schemes detailed exposition and simulation methodology," *arXiv preprint arXiv:2601.15004*, 2026.

[16] S. A. A. Hakeem and H. Kim, "Pq-v2x: A novel post-quantum cryptographic dataset for secure vehicular communications," *IEEE Internet of Things Journal*, 2025.

[17] Y. Ishai, E. Kushilevitz, M. Prabhakaran, A. Sahai, and C.-H. Yu, "Secure protocol transformations," in *Annual International Cryptology Conference*. Springer, 2016, pp. 430–458.

[18] D. Dragomir, L. Gheorghe, S. Costea, and A. Radovici, "A survey on secure communication protocols for iot systems," in *2016 international workshop on Secure Internet of Things (SIoT)*. IEEE, 2016, pp. 47–62.

[19] S. M. Mohammed, A. Al-Barrak, and N. T. Mahmood, "Enabling technologies for ultra-low latency and high-reliability communication in 6g networks." *Ingenierie des Systemes d'Information*, vol. 29, no. 3, 2024.

[20] M. U. A. Siddiqui, H. Abumarshoud, L. Bariah, S. Muhaidat, M. A. Imran, and L. Mohjazi, "Urllc in beyond 5g and 6g networks: An interference management perspective," *IEEe Access*, vol. 11, pp. 54 639–54 663, 2023.

[21] B. Akande, "The impact of quantum computing on encryption: How quantum computers can break current encryption methods, such as rsa and ecc, and what this means for data security," 2025.

[22] I. A. Abdulrahman, C. Anadozie, J. Alebiosu, G. E. Egbedion, G. T. Ayodele, and O. Eniola, "Quantum computing and its impact on cryptography: The future of secure communications and post-quantum cryptography."

[23] L. Chen, D. Moody, R. Perlner, and D. Smith-Tone, "Nist post-quantum cryptography: A preliminary performance evaluation," *NIST Internal Report*, 2021.

[24] D. Bhatt, N. Bhardwaj, and O. Pal, "Pqc for 5g and beyond: Foundations, challenges, and future directions," *Challenges, and Future Directions*.

[25] Y. Huang, M. Huang, Z. Lei, and J. Wu, "A pure hardware implementation of crystals-kyber pqc algorithm through resource reuse," *IEICE Electronics Express*, vol. 17, no. 17, pp. 20 200 234–20 200 234, 2020.

[26] N. Lohmiller, S. Kaniewski, M. Menth, and T. Heer, "A survey of post-quantum cryptography migration in vehicles," *IEEE Access*, 2025.

[27] H. Nguyen, S. Huda, Y. Nogami, and T. T. Nguyen, "Security in post-quantum era: A comprehensive survey on lattice-based algorithms," *IEEE Access*, 2025.

[28] M. Abbasi, F. Cardoso, P. Váz, J. Silva, and P. Martins, "A practical performance benchmark of post-quantum cryptography across heterogeneous computing environments," *Cryptography*, vol. 9, no. 2, p. 32, 2025.

[29] M. Rushad, A. Nambiar, and B. Chandavarkar, "Resource-aware cryptography: an analysis of lightweight cryptographic primitives," *SN Computer Science*, vol. 3, no. 1, p. 98, 2022.

[30] S. Kaganurmath and N. Cholli, "Enabling robust security in mqtt-based iot networks with dynamic resource-aware key sharing," *Procedia Computer Science*, vol. 252, pp. 633–642, 2025.

[31] K. Mahmood, S. Khan, M. Abdelhaq, M. U. Hassan, M. Uddin, R. Alsaqour, K. A. Awan, and M. A. Alsoufi, "Adaptive resource aware and privacy preserving federated edge learning framework for real time internet of medical things applications," *Scientific Reports*, vol. 15, no. 1, p. 36468, 2025.

[32] T. Jager, J. Schwenk, and J. Somorovsky, "On the security of tls 1.3 and quic against weaknesses in pkcs# 1 v1. 5 encryption," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 2015, pp. 1185–1196.

[33] S. Cook, B. Mathieu, P. Truong, and I. Hamchaoui, "Quic: Better for what and for whom?" in *2017 IEEE International Conference on Communications (ICC)*. IEEE, 2017, pp. 1–6.

[34] L. Crabbe and M. A. Post, "The effect of a rating downgrade on outstanding commercial paper," *The Journal of Finance*, vol. 49, no. 1, pp. 39–56, 1994.

[35] Y. Joarder and C. Fung, "Exploring quic security and privacy: A comprehensive survey on quic security and privacy vulnerabilities, threats, attacks and future research directions," *IEEE Transactions on Network and Service Management*, 2024.

[36] Y. A. Ergu, V.-L. Nguyen, R.-H. Hwang, Y.-D. Lin, C.-Y. Cho, H.-K. Yang, H. Shin, and T. Q. Duong, "Efficient adversarial attacks against drl-based resource allocation in intelligent o-ran for v2x," *IEEE Transactions on Vehicular Technology*, 2024.

[37] B. Flowers, R. M. Buehrer, and W. C. Headley, "Evaluating adversarial evasion attacks in the context of wireless communications," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1102–1113, 2019.

[38] P. Ravi, S. Bhasin, S. S. Roy, and A. Chattopadhyay, "On exploiting message leakage in (few) nist pqc candidates for practical message recovery attacks," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 684–699, 2021.

[39] S. Thangam and S. S. Chakkaravarthy, "An edge enabled region-oriented dag-based distributed ledger system for secure v2x communication," *KSII Transactions on Internet and Information Systems (TIIS)*, vol. 18, no. 8, pp. 2253–2280, 2024.

[40] P. Arcaini, E. Riccobene, and P. Scandurra, "Modeling and analyzing mape-k feedback loops for self-adaptation," in *2015 IEEE/ACM 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. IEEE, 2015, pp. 13–23.

APPENDIX

A. Proof of Decision Stability

Proof. Let $L_t(a)$ denote the weighted loss of algorithm a at time t, and let $\hat{L}_{t+1}(a)$ be the predicted loss computed from $\hat{X}_{t+1}$. By assumption (A1), each cost term is L-Lipschitz in context:

$$|L(a, X) - L(a, X')| \leq K\|X - X'\|, \tag{8}$$

for some constant K determined by the combined Lipschitz constants of all components of L. Prediction error satisfies $\|\hat{X}_{t+1} - X_{t+1}\| \leq \varepsilon$ by (A3). Thus, $|\hat{L}_{t+1}(a) - L_{t+1}(a)| \leq K\varepsilon$.

Let $\Delta_{\min}$ be the minimum loss separation between any two PQC algorithms:

$$\Delta_{\min} = \min_{a \neq b} |L_{t+1}(a) - L_{t+1}(b)|. \tag{9}$$

If $K\varepsilon < \Delta_{\min}$, then the perturbation induced by prediction error is insufficient to change the ordering of $L_{t+1}(a)$ across all $a \in \mathcal{A}$. Therefore, the identity of the minimizer is invariant:

$$\arg\min_a \hat{L}_{t+1}(a) \;=\; \arg\min_a L_{t+1}(a). \tag{10}$$

Since APMOEA selects based on the predicted loss, it returns the same algorithm as the oracle using the true context. Hence the decision remains stable for all t within any interval satisfying the loss-gap condition, proving the claim. □

B. Proof of Monotonic Upgrade Security

Proof. Let v_t denote the currently active PQC version and let $v_{t+1} > v_t$ be the version proposed by the optimizer. Each transition message includes an authenticated version counter ctr, a context hash $H(X_t)$, a freshness nonce N_t, and a digital signature generated within trusted hardware.

We show that no probabilistic polynomial-time (PPT) adversary can induce a downgrade or desynchronization. A replayed message necessarily contains an old nonce $N_{t'}$ and an outdated version counter $v_{t'} < v_{t+1}$, which is rejected by freshness checks. If the adversary modifies $(v_{t'}, N_{t'})$ to any (v, N), signature verification fails since both values are covered by the signed payload. To force a downgrade, the adversary must produce a valid message with $v < v_t$. However, modifying the version counter invalidates the signature, and forging a valid signature is infeasible under the assumed security of the PQC signature scheme. Both endpoints also enforce a monotonic version check, $v_{\text{received}} \geq v_t$, rejecting any message with $v < v_t$ regardless of its contents. The protocol further enforces a two-phase transition: an endpoint switches to v_{t+1} only after receiving a valid, fresh, signed acknowledgment from its peer. Any dropped, delayed, or altered message results in a mismatch and causes rollback to v_t, preventing unilateral advancement. Since no adversarial strategy can produce a valid transcript leading to $v < v_{t+1}$ or to inconsistent endpoint states, all accepted transitions are strictly monotonic and synchronized. □

C. Proof of Latency Boundedness

Proof. Let $T_{\text{lat}}(a, X_t)$ be the end-to-end latency under context X_t. Assumption (A2) states that vehicular dynamics satisfy: $\|X_{t+1} - X_t\| \leq \delta$. Each component of T_{lat} is continuous in the context variables (encoding time, decoding time, network delay, and SNR-dependent effects). Thus T_{lat} is itself Lipschitz: $|T_{\text{lat}}(a, X_{t+1}) - T_{\text{lat}}(a, X_t)| \;\leq\; C\delta$, for some $C > 0$. Let, $a_t = \arg\min_{a \in \mathcal{A}} L_t(a)$ be the algorithm selected by APMOEA which guarantees that a_t lies on, or near the Pareto front of the feasible set at time t. Then for any $a \in \mathcal{A}$, $T_{\text{lat}}(a_t, X_t) \;\leq\; \max_{a' \in \mathcal{A}} T_{\text{lat}}(a', X_t)$ by definition of the maximum. At time $t+1$, we have$T_{\text{lat}}(a_t, X_{t+1}) \leq T_{\text{lat}}(a_t, X_t) + C\delta$, Thus,

$$T_{\text{lat}}(a_t, X_t) \;\leq\; \max_{a' \in \mathcal{A}} T_{\text{lat}}(a', X_t) + O(\delta), \tag{11}$$

which proves that latency growth remains bounded whenever context drift is bounded. In vehicular environments, δ corresponds to changes in mobility, fading, and load over one optimization interval (20–50 ms), which are small. Hence PQC latency remains URLLC-compliant. □

D. Robustness Under Small Prediction Error

Lemma A.1 (Robustness Under Small Prediction Error). *If the prediction error satisfies the stability condition $K\varepsilon < \Delta_{\min}$, then APMOEA selects $a_t = a^\star$.*

APMOEA under admissible prediction noise matches the oracle's choice for optimal PQC, with secure migration and detection mechanisms as outlined in Algorithm 2.

Proof. Let, $a^\star = \arg\min_{a \in \mathcal{A}} L_{t+1}(a)$ be the oracle-optimal PQC algorithm under the true next context X_{t+1}. APMOEA computes: $a_t = \arg\min_a \hat{L}_{t+1}(a)$, where $\hat{L}_{t+1}(a)$ is computed using the predicted context $\hat{X}_{t+1}$. By (A3), $\|\hat{X}_{t+1} - X_{t+1}\| \leq \varepsilon$. By Lipschitz continuity, $|\hat{L}_{t+1}(a) - L_{t+1}(a)| \leq K\varepsilon$. If $K\varepsilon < \Delta_{\min}$ (the stability condition), then the perturbation in losses is insufficient to change the identity of the minimizer. Thus,

$$\arg\min_a \hat{L}_{t+1}(a) \;=\; \arg\min_a L_{t+1}(a) \;=\; a^\star. \tag{12}$$

Therefore, APMOEA matches the oracle's decision whenever prediction errors are small relative to the loss separation. This guarantees correct PQC selection under admissible noise. □

Table III compares four NIST PQC families in terms of computational, communication, and robustness properties. Table IV shows that adaptive PQC selection improves URLLC latency and switching stability over static baselines. Figure 3 evaluates robustness via cost stability, selection behavior across NSGA-II, RL-only, and APMOEA, and runtime feasibility. Figure 4 confirms secure downgrade resistance and links higher prediction error to increased switching frequency.

E. Operational definitions used in results

End-to-end latency. Latency is defined as $T_{\text{e2e}}(a, X_t) = T_{\text{crypto}}(a, \text{CPU}_t) + T_{\text{net}}(\text{SNR}_t, \text{PER}_t)$. Cryptographic delay is taken from NIST reference implementations and automotive CPU profiles, while network delay follows an SNR/PER-based NR-V2X model averaged over 200 Monte Carlo runs. Results are evaluated against a 5–20 ms URLLC budget.

Cost under context manipulation. Honest cost reflects nominal context evolution, while attacked cost applies bounded perturbations to SNR and compute load to assess selection robustness without modifying cryptographic primitives.

Switching stability. Stability is measured by PQC switching frequency within a fixed window (e.g., switches per 60 s), where lower values indicate more reliable operation.

Downgrade resistance. Downgrade protection is enforced via monotonic version checks, context consistency, nonce freshness, and atomic confirmation, preventing downgrade and desynchronization attacks.