ConfidSPEC-V2X: A Quantum-Blockchain Intelligence for Mitigating Confidentiality Threats in Vehicle-to-Everything Networks

Collins Izuchukwu Okafor, Love Allen Chijioke Ahakonye, *Senior Member, IEEE*, Dong-Seong Kim, *Senior Member, IEEE*, and Jae-Min Lee, *Member, IEEE*

***Abstract*—Vehicular-to-everything (V2X) communications promise unprecedented safety and efficiency gains but remain vulnerable to confidentiality breaches such as eavesdropping, traffic analysis, and man-in-the-middle (MITM) attacks. We propose ConfidSPEC-V2X, a focused hybrid framework that integrates continuous-variable quantum key distribution (CV-QKD), a multiagent deep reinforcement learning (MADRL), and an Ethereum-based permissioned blockchain PureChain public-key infrastructure (PKI) to deliver information-theoretic secrecy, dynamic traffic obfuscation, and tamper-proof key management. In the quantum module, CV-QKD transceivers embedded in on-board units (OBUs) and roadside units (RSUs) establish symmetric keys resilient to passive interception and capable of immediate eavesdropping detection. The artificial intelligence (AI) module employs multiagent DRL agents at RSUs to learn optimal dummy-traffic injection policies that obfuscate real V2X message patterns against statistical inference. The blockchain module leverages PureChain smart contracts to register, rotate, and timestamp vehicle public keys, ensuring that any MITM attempt to forge or replay keys is invalidated. We implement and evaluate ConfidSPEC-V2X within an OMNeT++/Veins simulation under realistic urban mobility scenarios, measuring the quantum bit error rate (QBER), key generation throughput, obfuscation entropy, and key management latency. Results demonstrate that our framework achieves robust confidentiality protection with minimal performance overhead.**

***Index Terms*—Blockchain, eavesdropping, IoV man-in-the-middle (MITM), quantum key distribution (QKD), traffic analysis, vehicles and roadside infrastructure (V2I), vehicular-to-everything (V2X).**

Received 8 October 2025; revised 25 February 2026; accepted 16 April 2026. Date of publication 22 April 2026; date of current version 9 July 2026. This work was supported in part by the Innovative Human Resource Development for Local Intellectualization Program through the Institute for Information and Communications Technology Planning & Evaluation (IITP) Grant funded by Korea Government (Ministry of Science and ICT (MSIT), South Korea, 20%) under Grant IITP-2026-RS-2020-II201612; in part by the Priority Research Centers Program through the National Research Foundation of Korea (NRF) funded by Ministry of Education, Science and Technology (MEST) (20%) under Grant 2018R1A6A1A03024003; in part by MSIT through the Information Technology Research Center (ITRC) Support Program (20%) under Grant IITP-2026-RS-2024-00438430; in part by the Basic Science Research Program through NRF funded by the Ministry of Education (20%) under Grant RS-2025-25431637; and in part by the Regional Innovation System and Education (RISE)-(Regional Growth Innovation Laboratory) Program through Gyeongbuk RISE Center, funded by the Ministry of Education (MOE), Gyeongsangbuk-do, Republic of Korea (20%), under Grant 2026-rise-15-105. *(Corresponding author: Jae-Min Lee.)*

Collins Izuchukwu Okafor and Jae-Min Lee are with the Department of IT Convergence Engineering, Kumoh National Institute of Technology, Gumi 39177, South Korea (e-mail: collinsokafor@kumoh.ac.kr; ljmpaul@kumoh.ac.kr).

Love Allen Chijioke Ahakonye is with the ICT Convergence Research Center, Kumoh National Institute of Technology, Gumi 39177, South Korea (e-mail: loveahakonye@kumoh.ac.kr).

Dong-Seong Kim is with the Department of IT Convergence Engineering, NSLab Company Ltd., Kumoh National Institute of Technology, Gumi 39177, South Korea (e-mail: dskim@kumoh.ac.kr).

Digital Object Identifier 10.1109/JIOT.2026.3686364

I. Introduction

VEHICLE-TO-EVERYTHING (V2X) communication stands as a cornerstone technology for the next generation of intelligent transportation systems (ITSs), promising transformative improvements in road safety, traffic efficiency, and the progressive realization of autonomous driving capabilities [1]. By enabling real-time, low-latency information exchange between vehicles (V2V), vehicles and roadside infrastructure (V2I), vehicles and vulnerable road users (V2P), and vehicles and the wider network (V2N), V2X facilitates a rich ecosystem of cooperative awareness applications [2], [3]. These include, but are not limited to, collision avoidance warnings, hazardous location alerts, optimized traffic signal phasing, cooperative maneuvering, and advanced driver-assistance systems (ADAS) augmentation [4]. However, the inherent openness of the wireless medium and the broadcast nature of many V2X messages, such as basic safety messages (BSMs) or cooperative awareness messages (CAMs), create a significantly expanded attack surface [5]. This makes robust security and privacy measures imperative, with confidentiality emerging as a particularly paramount requirement to ensure user trust and system integrity [5], [6].

Ensuring the confidentiality of V2X communications is critical due to the sensitive nature of the data exchanged. Messages frequently contain precise spatio–temporal data (location, speed, and heading), temporary pseudonyms that could potentially be linked to real-world identities through sustained observation, granular driving behavior patterns (acceleration and braking habits), and even safety-critical control information in more advanced V2X use cases [6], [7]. Unauthorized access to, or inference from, this information can lead to severe consequences. These range from profound privacy violations, such as persistent tracking of individuals leading to physical stalking or unwanted profiling for commercial purposes [8], to enabling more sophisticated cyber–physical attacks that can directly compromise vehicle safety, disrupt traffic management, or undermine the overall integrity of the ITS [9]. Within this context, several primary confidentiality threats warrant significant concern in the V2X landscape.

First, eavesdropping involves the unauthorized passive interception of wireless V2X messages by an adversary. For instance, an attacker positioned within radio range could capture BSMs to reconstruct a vehicle's trajectory, identify frequently visited locations, or infer personal routines [9], [10]. In a V2I context, eavesdropping on communications between a vehicle and the charging infrastructure could reveal payment details or energy consumption patterns. The consequences extend beyond individual privacy; aggregated eavesdropped data could facilitate large-scale surveillance or corporate espionage (e.g., tracking commercial fleet movements). More critically, information gleaned from eavesdropping, such as session identifiers or cryptographic nonces, if poorly protected, might be used as a precursor to more active attacks [11].

Following this, even when V2X messages are encrypted, traffic analysis presents a potent threat, where adversaries infer significant sensitive information by meticulously analyzing communication metadata and traffic flow patterns [12], [13]. For example, an attacker observing the volume, frequency, and source-destination pseudonyms of encrypted messages between specific vehicles could deduce the formation or dissolution of platoons, identify vehicles frequently co-located or interacting (potentially deanonymizing participants over time through linkage attacks), or recognize abnormal communication bursts indicative of an emergency event or a specific V2X service activation [14]. Such inferences can compromise operational security, reveal social graphs, or provide contextual intelligence for targeting further attacks, all without requiring the decryption of the actual message content.

Furthermore, man-in-the-middle (MITM) attacks pose an active danger, where an adversary intercepts the communication channel between two legitimate V2X entities [e.g., two vehicles, or a vehicle and a roadside unit (RSU)], often by impersonating one or both parties. A common MITM scenario involves an attacker positioning themselves between a vehicle and an RSU to intercept, potentially modify, and then relay messages, making both legitimate parties believe they are communicating directly with each other [9]. For example, an attacker could impersonate an RSU to send false traffic light information (SPaT messages) or emergency warnings, causing vehicles to take dangerous actions. Conversely, they could impersonate a vehicle to inject false BSMs indicating a nonexistent hazard, leading to unnecessary traffic disruption [14]. MITM attacks can also target over-the-air (OTA) software update mechanisms, intercepting update packages to inject malware or prevent legitimate updates, severely compromising vehicle functionality and security [15]. Such attacks directly threaten the authenticity and confidentiality of exchanged data, leading to immediate safety risks and erosion of trust in the V2X system.

Standard cryptographic techniques, such as encryption using symmetric and asymmetric algorithms, form a foundational layer of defense. However, they present inherent limitations in comprehensively addressing these confidentiality threats [16]. While encryption protects message content, it does not inherently prevent sophisticated traffic analysis, as metadata remains observable. Traditional public-key infrastructures (PKIs), typically reliant on centralized certificate authorities (CAs) for managing cryptographic identities and authenticating V2X participants, can become single points of failure. If a CA is compromised, an attacker could issue fraudulent certificates, enabling widespread MITM attacks [17].

Furthermore, the advent of large-scale, fault-tolerant quantum computers poses an existential threat to currently deployed public-key cryptography, including RSA and elliptic curve cryptography (ECC), which underpin most secure communication protocols today [16], [18], [19]. Shor's algorithm, for instance, can efficiently break these systems, rendering encrypted data (past and future) vulnerable. While postquantum cryptography (PQC) aims to develop new classical cryptographic algorithms presumed to be resistant to quantum attacks and is undergoing standardization by NIST [20], PQC algorithms themselves can introduce new challenges. Many PQC candidates, particularly for key encapsulation and digital signatures, often involve significantly larger key sizes, larger signature sizes, and higher computational overhead compared to their ECC counterparts [6], [21]. For instance, IoT networks often use key sizes of 128–4096 bits, while PQC key sizes can range from thousands of kilobytes to megabytes [22]. These characteristics can be particularly problematic for resource-constrained V2X on-board units (OBUs) and RSUs, potentially impacting communication latency, bandwidth usage, and hardware costs. Thus, while PQC is crucial for securing data-at-rest and digital signatures in a postquantum world, relying solely on PQC for all aspects of key exchange for confidentiality may not be the optimal or most robust long-term strategy, primarily against passive eavesdroppers who can record now and decrypt later.

Addressing these multifaceted confidentiality challenges effectively necessitates a more sophisticated, resilient, and forward-looking approach that strategically combines the strengths of different technological paradigms. This article introduces ConfidSPEC-V2X, a novel security framework meticulously designed to bolster confidentiality in V2X networks through targeted mitigation strategies. The core philosophy of ConfidSPEC-V2X is the strategic and modular application of the most appropriate state-of-the-art technologies: PQC, advanced artificial intelligence (AI), and blockchain, to specifically counter each distinct confidentiality threat vector. Instead of a one-size-fits-all solution, ConfidSPEC-V2X advocates for employing the information-theoretic security of continuous-variable quantum key distribution (CV-QKD) to provide unparalleled, physics-based protection against eavesdropping by ensuring secure key exchange. Concurrently, it proposes utilizing the adaptive learning power of multiagent deep reinforcement learning (MADRL)-based AI agents, deployed at RSUs, to dynamically obfuscate V2X communication patterns and thereby thwart traffic analysis. To complete this triad of confidentiality measures, the framework leverages the decentralized trust, immutability, and Smart contract capabilities of an Ethereum-based blockchain, PureChain, to implement a robust PKI designed for secure key management, including rotation and timestamped validation, making it highly resilient against MITM attacks.

The main contributions of this article are as follows.

1) The proposal of the ConfidSPEC-V2X framework, detailing its targeted, multitechnology architecture for enhancing V2X confidentiality.
2) In-depth descriptions of the operational mechanics and V2X-specific integration of CV-QKD for eavesdropping prevention, MADRL for traffic analysis mitigation, and a PureChain-based PKI for countering MITM attacks.
3) A compelling rationale justifying the selection of each specific technology (quantum, AI, and blockchain) as the optimal countermeasure for its assigned confidentiality threat within the V2X context.
4) A preliminary discussion of the potential synergies between these solutions, alongside implementation challenges and performance considerations pertinent to the proposed framework.

This work aims to provide a robust blueprint for a next-generation security architecture focused explicitly on preserving the confidentiality of communications within the increasingly complex, interconnected, and vital V2X ecosystem.

The remainder of this article is structured as follows. Section II reviews related work on confidentiality threats and challenges in current V2X communication frameworks. Section III provides a conceptual overview and system model of the ConfidSPEC-V2X framework, delving into the detailed mechanisms for mitigating Eavesdropping, Traffic Analysis, and MITM attacks, respectively, using the chosen technologies. Section IV discusses the experimental results and performance considerations of the framework. Finally, Section V concludes the article, summarizing its contributions, impact, and promising future research directions.

II. Background Study and Related Works

This section analyzes the reviewed literature on confidentiality threats and mediation approaches in V2X communication, as presented by various researchers. Although recent work highlights the need for scalable, reliable, and communication-efficient designs to enable quantum-resistant protocols, unbreakable wide-area communications, and secure, scalable decentralized identity for Web3 [23], [24]. This review focuses on three critical confidentiality threats in V2X networks: eavesdropping, traffic analysis, and MITM attacks. For each threat, recent countermeasures are examined, their inherent limitations are identified, and a critical analysis is provided on how the ConfidSPEC-V2X framework, detailed in this work, offers potentially more robust and sustainable solutions.

A. Eavesdropping Attacks and Mitigation Approaches

Studies by Fowler et al. [25] proposed a practical V2I system using free-space optical quantum key distribution (FSO-QKD) for key delivery. This approach integrates a QKD link to continuously generate symmetric keys for V2I communication, replacing conventional PKI. The system includes a zero-trust authentication protocol, ensuring that spoofed or eavesdropped messages (e.g., fake V2I messages) are rejected. Free-space QKD approaches require line-of-sight optics, which are distance/rate-constrained, infrastructure-heavy, and come with a very high integration complexity. In this work, CV-QKD (more channel-tolerant) is used for keys, and blockchain hashes to authenticate identities, while combining quantum eavesdropper detection with ledger-based PKI.

Tang et al. [26] address eavesdropping at the physical layer by leveraging reconfigurable intelligent surfaces (RISs). They consider a hybrid V2V/V2I scenario where an external eavesdropper passively listens. They formulate a secrecy-rate maximization problem (the "physical layer secure transmission rate" or PLSTR) for the V2I base station, optimizing its precoding and the RIS reflection coefficients via alternating optimization (generalized Rayleigh quotient and semidefinite relaxation). Their core contribution is an algorithm that jointly configures multiple-input–multiple-output (MIMO) beamforming and RIS phases to nullify the eavesdropper's channel while preserving the intended V2I link. They validate via simulation that significant secrecy-rate gains are achievable and analyze how RIS parameters affect performance. However, their scheme assumes perfect channel state information (CSI), including the eavesdropper's channel, and a static or slowly varying topology. In practice, obtaining or tracking the eavesdropper CSI is difficult. The optimization is also complex and may not adapt quickly to mobility. In addition, RIS hardware must be deployed (with added cost and control signaling). ConfidSPEC-V2X addresses these limitations by bypassing reliance on perfect CSI: QKD intrinsically detects eavesdropping without needing an adversary model and provides information-theoretic secrecy independent of channel conditions.

B. Traffic Analysis Attack and Countermeasures

Traffic analysis attacks represent a sophisticated threat vector where adversaries analyze communication patterns and metadata to infer sensitive information without directly accessing message content. Recent research has focused on developing intelligent obfuscation techniques and privacy-preserving mechanisms to counter these analytical attacks.

Abdellah et al. [27] employ a bidirectional long short-term memory (BiLSTM) network to handle the temporal dependencies inherent in vehicular traffic patterns. The research established that intelligent traffic prediction mechanisms could be reversed to develop sophisticated obfuscation strategies that make traffic analysis significantly more challenging for potential adversaries. However, the approach faces limitations in its direct application to security contexts, as the same predictive capabilities that benefit legitimate traffic management could potentially be exploited by sophisticated adversaries for enhanced traffic analysis attacks.

Worae and Mastorakis [28] propose "hiding in plain sight," an IoT traffic obfuscation framework to thwart machine-learning traffic analysis. Unlike one-shot padding methods, they integrate six techniques, packet padding, XOR-based padding, shifting, constant-size padding, fragmentation, and randomized delays into a single pipeline. Their contribution is a comprehensive "camouflage" approach that systematically

blinds traffic classifiers by making packet features uniform or noisy. Injecting delays and padding inflates traffic volume and delays real communication. In constrained IoT links, this can hurt the quality of service. Moreover, their policy is static per-device (determined offline), so it may be suboptimal in dynamic networks. Finally, adversaries could attempt to filter out obvious decoy patterns (e.g., constant padding overhead). ConfidSPEC-V2X could improve on this by using deep reinforcement learning (DRL)-based adaptive obfuscation. Rather than fixed rules, a DRL agent could learn the minimal obfuscation needed in real-time to confuse an online traffic classifier, balancing privacy versus latency.

C. MITM Attack and Prevention Mechanisms

MITM attacks pose significant threats to V2X communication integrity by enabling adversaries to intercept, modify, and relay messages between legitimate parties. Recent research has concentrated on developing robust authentication mechanisms and decentralized trust models to prevent such attacks.

Shewajo et al. [17] propose a blockchain-PKI architecture for 5G V2X. They note that standard PKI (as in ETSI ITS) is centralized and noninteroperable, leading to single points of failure and difficulty roaming across domains. Their solution is a permissioned blockchain (Hyperledger Fabric) that spans multiple operators and the 5G Core. Vehicle and RSU certificates are stored (or referenced) in the chain, enabling any party to verify authenticity against a decentralized ledger. This yields a distributed trust model, which resists single-point compromises and eases cross-domain authentication. The limitation is added complexity and latency. Blockchain write/read operations (even on Fabric) take time and consume bandwidth. Managing on-chain revocations, key updates, and access control policies is a nontrivial task. ConfidSPEC-V2X already adopts a similar blockchain-PKI concept for authentication, aligning well with Shewajo et al.'s approach. ConfidSPEC would further benefit by integrating CV-QKD keys as the root of trust: the blockchain can log QKD-generated public keys or sessions, linking quantum-protected keys to vehicle identities. In this way, ConfidSPEC retains all the blockchain-based interoperability of Shewajo et al. [17], but adds an extra layer of key security via QKD (removing PKI's single-failure risk).

Farran et al. [29] proposed a blockchain-based V2X communication system that leverages distributed public keystore platforms built on Ethereum blockchain technology. Their approach replaces traditional PKI with a decentralized trust model that eliminates single points of failure while providing tamper-resistant key management capabilities. The system implements smart contracts for automated key storage, retrieval, and validation processes, ensuring that certificate management operations are transparent and verifiable. The primary limitations of blockchain-based approaches relate to scalability and transaction costs associated with the operation of distributed ledgers. The research indicated that certificate registration operations require significant transaction fees (approximately 0.00548 ETH per registration), which could become prohibitive for large-scale deployments. In addition, the dependency on blockchain network availability and performance could introduce vulnerabilities in scenarios where network connectivity is intermittent or compromised. The use of PureChain, a permissioned blockchain, in ConfidSPEC-V2X with smart contracts for automated, cost-effective key management.

This comprehensive analysis demonstrates that while significant progress has been made in addressing individual V2X confidentiality threats, the lack of integrated solutions represents a critical research gap. The ConfidSPEC-V2X framework's multitechnology approach provides a strategic response to these limitations by combining the strengths of quantum cryptography, AI, and blockchain technology in a coordinated security architecture designed explicitly for V2X environments.

III. System Model

The ConfidSPEC-V2X framework is designed as a three-modular architecture framework with each of the modules addressing a V2X confidentiality threat, such as eavesdropping, traffic analysis, and MITM threats.

A. Mitigating Eavesdropping With QKD (Module I)

To address the fundamental threat of eavesdropping in the V2X wireless channel, Module I of the ConfidSPEC-V2X framework employs CV-QKD. Quantum key distribution (QKD) protocols are broadly categorized into two main families: discrete-variable (DV) QKD and continuous-variable (CV) QKD. While both aim to establish a secure cryptographic key, they differ fundamentally in how they encode, transmit, and detect quantum information [30], [31]. DV-QKD protocols, such as BB84, rely on encoding information in discrete quantum states of single photons (e.g., polarization or phase) [32]. Security arises from the no-cloning theorem: any attempt by an eavesdropper to measure or copy these quantum states introduces detectable errors in the key, typically quantified via the quantum bit error rate (QBER) [33]. In contrast, CV-QKD encodes information into the continuous quadratures (amplitude and phase) of coherent states. These are measured with homodyne or heterodyne detectors, which can leverage standard telecommunication components instead of single-photon detectors [34], [35]. While DV-QKD generally provides robustness over long distances, it suffers from low key generation rates and expensive hardware requirements. CV-QKD, on the other hand, achieves higher secret key rates and better integration with classical optical infrastructure, though it is more sensitive to excess channel noise [36]. Unlike classical cryptographic methods that rely on computational complexity or DV-QKD, like BB84 that relies on a single photon for the transmission of quantum information, CV-QKD provides information-theoretic security for key exchange, guaranteed by the laws of quantum mechanics, with higher key rates and more robust to channel losses in the low phase noise regime but high thermal noise regime [37]. Any attempt by an eavesdropper, Eve, to intercept and measure the quantum states inevitably introduces detectable disturbances, revealing her presence. This section outlines the system model

and protocol steps for implementing a prepare-and-measure CV-QKD scheme using Gaussian-modulated coherent states (GMCSs), which is particularly well-suited for integration with existing telecommunication technologies.

1) System Model and Protocol Participants: The system model consists of two legitimate V2X entities, conventionally named Alice and Bob, who wish to establish a secret symmetric key.

1) *Alice (Transmitter):* An OBU or RSU that prepares and sends a sequence of quantum states (faint laser pulses) to Bob.
2) *Bob (Receiver):* The corresponding OBU or RSU that receives and measures the quantum states sent by Alice.
3) *Quantum Channel:* The physical medium between Alice and Bob (e.g., free-space optical link or fiber optic cable) through which the quantum states are transmitted. It is characterized by loss and noise.
4) *Classical Public Channel:* An authenticated classical communication channel (which can be the standard V2X network itself, secured by a preexisting mechanism like digital signatures from Module III) used for postprocessing steps. Eve is assumed to have full access to this channel, but cannot tamper with it due to the authentication.

The protocol proceeds in two main phases: the quantum transmission phase, where quantum states are exchanged, and the classical postprocessing phase, where a secure key is distilled from the measurement outcomes.

2) Quantum Transmission Phase: State Preparation and Measurement: The core of the quantum exchange is based on the GMCS protocol.

Step 1 [State Preparation (Alice)]: Alice generates a sequence of N coherent states $|\alpha_k\rangle$. For each state $k = 1, \ldots, N$, she generates a pair of random numbers, $(x_A^{(k)}, p_A^{(k)})$, drawn from a zero-mean Gaussian distribution with a modulation variance of σ_A^2. These numbers represent the displacements of the state in the phase space along the X (amplitude) and P (phase) quadratures. In our model, V_A represents the total added modulation variance, so the variance per quadrature is $V_{A,\text{mod\_quad}} = V_A/2$.

The complex amplitude α_k of the coherent state is given by

$$\alpha_k = \frac{x_A^{(k)} + i p_A^{(k)}}{\sqrt{2}}. \tag{1}$$

The total variance of Alice's prepared signal quadrature, including the fundamental vacuum noise N_0, is $V_{\text{Alice, quad}} = V_{A,\text{mod\_quad}} + N_0$. In our simulation, using the convention where total vacuum noise is normalized to 1 ($V_{\text{shot}} = 1$), the vacuum noise per quadrature is $N_0 = 1/2$.

Step 2 (Quantum Channel Transmission): The prepared states are transmitted through the quantum channel, which inevitably introduces loss (transmittance T) and excess noise (ϵ). The relationship between the input state from Alice and the state received by Bob can be modeled linearly for each quadrature

$$\hat{x}_B = \sqrt{T}\hat{x}_A + \hat{N}_x \quad \text{and} \quad \hat{p}_B = \sqrt{T}\hat{p}_A + \hat{N}_p \tag{2}$$

where $\hat{N}_x$ and $\hat{N}_p$ are the noise terms. The total noise variance at Bob's side, referred to as the channel input, is modeled as $\chi_{\text{tot}} = (1/T - 1)N_0 + \epsilon$. The excess noise ϵ quantifies all noise sources beyond the fundamental vacuum noise introduced by loss, including thermal noise, detector noise, and any disturbance caused by Eve.

Step 3 [Homodyne Detection (Bob)]: For each incoming state, Bob randomly chooses to measure either the X-quadrature or the P-quadrature using a homodyne detector. His measurement outcome for the kth state will be a real number $m_B^{(k)}$ corresponding to the measured quadrature value, which includes the effects of channel loss and noise. Based on the validated behavior of our simulation, Bob's measurement outcomes, m_B, are scaled such that their variance is consistent with the physical quadrature variance model, where $N_0 = 1/2$ for vacuum.

3) Classical Postprocessing Phase: After transmitting all N states, Alice and Bob use the authenticated public channel to distill a key.

Step 4 (Data Sifting): Bob announces the sequence of quadratures he measured for each state. Alice and Bob discard all data where Bob's measurement basis does not correspond to the quadrature data Alice intended to use for that signal (in some protocols), resulting in a smaller set of correlated raw data for Alice (X_A) and Bob (X_B).

Step 5 (Parameter Estimation): Alice and Bob publicly reveal a random subset of their sifted raw data to estimate the channel parameters T and ϵ. This step is critical for security.

1) *Transmittance Estimation (T_{est}):* T is estimated from the linear relationship between Alice's modulation values and Bob's measurement outcomes. By performing a linear regression of Bob's data (Y) on Alice's data (X), we obtain a model $Y_i = mX_i + c$. The slope m corresponds to $\sqrt{T}$. Therefore,
$$T_{est} = \left(\text{slope from regression}\right)^2. \tag{3}$$
2) *Excess Noise Estimation (ϵ_{est}):* The excess noise is calculated from the variance of the residuals of the linear regression. The variance of the residuals, $V_{\text{noise\_out}}$, represents the noise Bob observes that is uncorrelated with Alice's modulation. This noise is composed of the fundamental vacuum noise N_0 and the channel's excess noise, scaled by transmittance. The input-referred excess noise ϵ (in shot noise units, SNUs) is estimated as
$$\epsilon_{est} = \frac{V_{\text{noise\_out}} - N_0}{T_{est} \cdot N_0}. \tag{4}$$

If T_{est} is significantly lower or ϵ_{est} is significantly higher than expected for the V2X link, Alice and Bob must assume the presence of an eavesdropper and abort the protocol.

Step 6 (Information Reconciliation and QBER): Due to channel noise, Bob's sifted data still contains errors relative to Alice's. They perform information reconciliation (error correction), typically using efficient codes like low-density parity-check (LDPC) codes over the public channel. The efficiency of this process is denoted by $\beta \in [0, 1]$. Before reconciliation, the QBER is calculated by comparing a subset of their discretized data to quantify the raw error level.

Step 7 (Privacy Amplification and SKR): Eve may have gained partial information about the key by observing the channel and the reconciliation process. To eliminate this information, Alice and Bob apply a universal hash function (e.g., SHA-256) to their error-corrected key string to distill a shorter, but information-theoretically secure, final key.

The length of this final key is determined by the secure key rate (SKR), which quantifies the achievable rate of secret bits per transmitted signal. For reverse reconciliation (where Bob's data is corrected to match Alice's), a lower bound on the SKR against collective Gaussian attacks is given by

$$K = f_{\text{sample}} \cdot \left[\beta I(A:B) - \chi(B:E)\right] \tag{5}$$

where $f_{\text{sample}} = (N-n)/N$ is the fraction of signals used for key generation after n signals were used for parameter estimation, β is the reconciliation efficiency, $I(A:B)$ is the Shannon mutual information between Alice and Bob, and $\chi(B:E)$ is the Holevo quantity, which bounds the information Eve can obtain on Bob's data. These quantities are calculated as

$$I(A:B) = \frac{1}{2}\log_2(1+\text{SNR}) \tag{6}$$

$$\text{SNR} = physicalprinciplesguaranteeitsconfidentiality \tag{7}$$

$$\chi(B:E) = \log_2(1+V_E) \quad \text{where} \quad V_E = \frac{\epsilon}{1-T} \tag{8}$$

where V_A and V_{shot} are the total modulation and shot noise variances, respectively. Upon successful completion, this module provides a fresh, secret symmetric key to a key management interface (KMI), which can then be used by higher level V2X security protocols (e.g., for AES encryption of BSMs) with the assurance that physical principles guarantee its confidentiality.

B. AI-Driven Traffic Analysis Mitigation Methodology (Module II)

The primary objective of this module is to thwart traffic analysis attacks by dynamically altering V2X communication patterns. An adversary performing traffic analysis does not need to decrypt messages; instead, they analyze metadata such as message frequency, timing, and size to infer sensitive information (e.g., platoon formation, vehicle-to-infrastructure service access). To counter this, we deploy an intelligent agent at each RSU tasked with injecting dummy traffic to obfuscate these patterns.

Unlike a standard Markov decision process (MDP), where the full system state is visible, our problem is modeled as a partially observable MDP (POMDP). Here, the RSU agent only perceives partial and noisy observations o_t generated from the hidden state s_t via the observation model $\mathcal{O}(o_t|s_t)$, reflecting its limited and uncertain view of the network traffic. The key components are as follows.

State (s_t): At any timestep t, the state is a vector of features representing the observed traffic pattern within the RSU's range. This includes metrics such as: average message frequency, interpacket delay distribution statistics (mean, variance), data volume per unit time, and source and destination pseudonym activity.

Action (a_t): The action is a continuous vector that defines the dummy traffic to be injected. This vector specifies parameters such as the rate of dummy packet injection, the size distribution of the dummy packets, and the timing of the injections.

Reward (R_t): The agent receives a reward at each timestep based on its success in obfuscating the traffic. The goal is to make the combined traffic (real + dummy) statistically indistinguishable from random noise. This is achieved by rewarding the agent for maximizing the entropy of the observed traffic pattern. A common way to formulate this is to have the agent minimize the Kullback–Leibler (KL) divergence between the observed traffic distribution, P_{obs}, and a target uniform distribution, U. The reward is thus the negative of this divergence

$$R_t = -D_{KL}(P_{obs}(s_t)\|U) \tag{9}$$

where D_{KL} means KL divergence. A lower KL divergence signifies that the observed traffic pattern is closer to random noise, resulting in a higher (less negative) reward.

Observation (o_t): At each timestep, the RSU agent does not have full knowledge of the global network state. Instead, it perceives a local observation vector o_t derived from partial traffic features such as interarrival times of packets, short-term message counts, and pseudonym usage within its coverage. These observations are noisy and incomplete, reflecting the limited sensing scope of the RSU.

Observation Model ($\mathcal{O}$): The observation model $\mathcal{O}(o_t|s_t)$ defines the probability of receiving observation o_t given the underlying hidden state s_t. In our setting, the model captures how the true traffic distribution within the RSU's range is mapped to the sampled metadata features. For instance, bursty state transitions may only be partially visible through sampled packet interarrival times, while hidden correlations (e.g., vehicle platoon dynamics) may remain unobserved. This probabilistic mapping ensures that the POMDP formulation correctly models the uncertainty faced by RSU agents.

1) PPO Algorithm: To solve this POMDP, we employ the proximal policy optimization (PPO) algorithm. PPO is a state-of-the-art policy gradient method chosen for its remarkable stability, reliability, and sample efficiency, which are critical for learning in dynamic V2X network environments. The core innovation of PPO is that it constrains the policy updates at each training step, preventing the new policy from deviating too far from the old one. This is achieved by optimizing a *clipped* surrogate objective function, which discourages significant, potentially destabilizing updates and prevents the kind of catastrophic performance collapse observed in earlier, less stable algorithms.

The PPO algorithm is built upon the actor–critic architecture, which consists of two main neural networks as follows.

1) *Actor (Policy Network):* This network, with parameters θ, learns the policy, $\pi_\theta(a_t|o_t)$, which maps a given observation o_t [sampled from the underlying hidden state s_t via the observation model $\mathcal{O}(o_t|s_t)$] to a probability distribution over possible actions a_t. For continuous action spaces, this is typically a Gaussian distribution.
2) *Critic (Value Network):* This network, with parameters ϕ, learns the value function, $V_\phi(o_t)$, which estimates

the expected total future reward (return) from a given observation o_t.

The training process involves optimizing the following components.

Advantage Function: The advantage function, $A(s_t, a_t)$, quantifies how much better a specific action a_t is compared to the average action that would be taken in state s_t. We use generalized advantage estimation (GAE) to calculate a reliable and stable estimate of the advantage.

First, we define the temporal difference (TD) error, δ_t

$$\delta_t = r_t + \gamma V_\phi\left(s_{t+1}\right) - V_\phi\left(s_t\right) \tag{10}$$

where r_t is the reward at time t and $\gamma \in [0, 1]$ is the discount factor.

The GAE advantage, $\hat{A}_t$, is then the exponentially weighted average of these TD errors

$$\hat{A}_t = \sum_{l=0}^{\infty} (\gamma\lambda)^l \delta_{t+l} \tag{11}$$

where $\lambda \in [0, 1]$ is the GAE parameter that controls the bias-variance tradeoff.

Clipped Surrogate Objective (Policy Loss): This is the central component of PPO. The objective is to maximize the expected advantage, but with a constraint that prevents the new policy π_θ from moving too far from the policy used to collect the data, $\pi_{\theta_{old}}$.

Let $r_t(\theta)$ be the probability ratio between the new and old policies

$$r_t\left(\theta\right) = \frac{\pi_\theta\left(a_t|s_t\right)}{\pi_{\theta_{old}}\left(a_t|s_t\right)}. \tag{12}$$

The clipped objective function, $L^{CLIP}(\theta)$, is then

$$L^{CLIP}\left(\theta\right) = \hat{\mathbb{E}}_t\left[\min\left(r_t\left(\theta\right)\hat{A}_t, \text{clip}\left(r_t\left(\theta\right), 1-\epsilon, 1+\epsilon\right)\hat{A}_t\right)\right] \tag{13}$$

where ϵ is a small hyperparameter (e.g., 0.2) that defines the clipping range. The "clip" function constrains the ratio $r_t(\theta)$ to the interval $[1 - \epsilon, 1 + \epsilon]$. The "min" operator ensures that the final objective is a lower bound (a pessimistic estimate) of the unclipped objective, which further discourages overly large policy updates.

Value Function Loss: The critic network is trained by minimizing the mean squared error (mse) between its value estimates $V_\phi(s_t)$ and the actual returns observed during the episode, V_t^{target} [often calculated as the sum of the advantages and the old value estimates, $\hat{A}_t + V_{\phi_{old}}(s_t)$)]

$$L^{VF}\left(\phi\right) = \hat{\mathbb{E}}_t\left[\left(V_\phi\left(s_t\right) - V_t^{target}\right)^2\right]. \tag{14}$$

Entropy Bonus: To encourage exploration and prevent the policy from collapsing to a suboptimal deterministic strategy too early, an entropy bonus is added to the loss function. The entropy S of the policy is given by

$$S\left[\pi_\theta\right]\left(s_t\right) = H\left(\pi_\theta\left(\cdot|s_t\right)\right). \tag{15}$$

Total Loss Function: The final loss function for PPO combines the policy loss, the value function loss, and the entropy bonus

$$L\left(\theta, \phi\right) = L^{CLIP}\left(\theta\right) - c_1 L^{VF}\left(\phi\right) + c_2 S\left[\pi_\theta\right]\left(s_t\right) \tag{16}$$

where c_1 and c_2 are coefficients that weight the importance of the value loss and the entropy bonus, respectively. This combined loss is then optimized using stochastic gradient descent.

Algorithm 1 PPO With GAE

Input: learning rates, clipparam ϵ, value loss coeff. c_1, entropy coeff. c_2, discount γ, GAE param. λ, batch size N, epochs K
Initialize policy parameters θ and value parameters ϕ
Initialize Actor network π_θ: input dim = $|o_t|$, hidden layers $[64, 64]$, output dim = $|a_t|$ (Gaussian head with mean and log-std)
Initialize Critic network V_ϕ: input dim = $|o_t|$, hidden layers $[64, 64]$, output dim = 1
Initialize Adam optimizers for θ, ϕ
for iteration = $1, 2, \ldots$ **do**
  clear storage
  for step = $1, \ldots, N$ **do**
    observe partial observation $o_t \sim \mathcal{O}(o_t|s_t)$
    sample $a_t \sim \pi_\theta(\cdot \mid o_t)$
    execute a_t, observe r_t, next observation o_{t+1},
    done flag d_t store

$$\left(o_t, a_t, r_t, d_t, \log \pi_\theta(a_t \mid o_t), V_\phi(o_t)\right)$$

  end for
  — compute GAE advantages —
  $\hat{A}_N = 0$
  for $t = N-1, \ldots, 1$ **downto do**
    $\delta_t \leftarrow r_t + \gamma\, V_\phi(o_{t+1})\,(1 - d_t) - V_\phi(o_t)$
    $\hat{A}_t \leftarrow \delta_t + \gamma\lambda\,(1 - d_t)\,\hat{A}_{t+1}$
  end for
  $R_t \leftarrow \hat{A}_t + V_\phi(o_t)$ (for all stored t)
  — perform K epochs of optimization —
  for epoch = $1, \ldots, K$ **do**
    for all minibatch $\left(o_b, a_b, \log \pi_{\theta_{\text{old}}}(a_b \mid o_b),\right.$
    $\left.\hat{A}_b, R_b\right)$ **do**
    $r(\theta) \leftarrow \dfrac{\pi_\theta(a_b \mid o_b)}{\pi_{\theta_{\text{old}}}(a_b \mid o_b)}$
    $L^{\text{CLIP}} \leftarrow \min\left(r(\theta)\,\hat{A}_b,\ \text{clip}(r(\theta), 1-\epsilon,\right.$
    $\left.1+\epsilon)\,\hat{A}_b\right)$
    $L^{\text{VF}} \leftarrow \left(V_\phi(o_b) - R_b\right)^2$
    $S \leftarrow \text{entropy}\left[\pi_\theta(\cdot \mid o_b)\right]$
    $L \leftarrow -\left(L^{\text{CLIP}} - c_1\, L^{\text{VF}} + c_2\, S\right)$
    update θ, ϕ by descending ∇L
    end for
  end for
end for

The complete training process for PPO is shown in Algorithm 1 below as follows.

C. Blockchain-Based PKI for MITM Mitigation (Module III)

To counter MITM attacks, Module III implements a decentralized PKI on a permissioned Ethereum-style blockchain (PureChain). Vehicles and RSUs become both PKI clients and blockchain peers, registering, rotating, and revoking their public keys via smart-contract calls. The key idea is that any

attempt by an adversary to inject a forged or replayed key can be detected (due to an invalid signature) or outrun by the next honest key-rotation transaction.

A MITM attack is one of the most severe threats to the integrity of V2X communication. In this scenario, an adversary, Eve, positions herself between two communicating entities (e.g., Vehicle A and Vehicle B, or a Vehicle and an RSU) and transparently relays messages between them. By controlling the communication channel, Eve can eavesdrop, inject malicious messages, or modify legitimate messages, all while making the two parties believe they are communicating directly and securely.

Traditional PKIs rely on a centralized CA to issue, manage, and revoke digital certificates. These certificates bind a public key to a specific identity. While effective, this centralized model presents a critical single point of failure. If the CA is compromised, an attacker can do the following.

1) Issue fraudulent certificates for malicious entities, allowing them to impersonate legitimate vehicles or infrastructure.
2) Illegitimately revoke the certificates of valid vehicles, causing a denial-of-service attack.
3) Delay the publication of certificate revocation lists (CRLs), creating a window of vulnerability where compromised keys can still be trusted.

These vulnerabilities make a centralized PKI an attractive target for attackers seeking to undermine the entire V2X network.

1) System Architecture: A Blockchain-Based PKI: Our proposed solution leverages the inherent security properties of blockchain technology to create a robust and resilient PKI.

Components:

1) *PureChain:* A permissioned Ethereum-based blockchain network. Only authorized entities can operate as validating nodes
2) *KeyManagement Contract:* On-chain CA that registers, rotates, and revokes V2X public keys.
3) *OBUs/RSUs:* Vehicles and RSUs, each with a unique ID and key pair, use PureChain to publish or fetch keys.

2) Core Operations and Smart Contract Logic: The "*KeyManagement*" smart contract is the core of this module. It maintains a secure, append-only ledger of all vehicle and RSU public keys.

3) Key Registration: When a new vehicle is manufactured, it is assigned a unique, immutable vehicle identity (VID). The manufacturer, acting as a trusted authority, initiates a transaction on PureChain to register the vehicle. The vehicle's OBU generates its initial key pair (sk_V, pk_V). The manufacturer creates a transaction to call the "registerVehicle" function of the "KeyManagement" smart contract. The transaction includes the vehicle's VID and its public key pk_V. Upon successful validation by the network, the smart contract creates a new record, mapping the VID to its pk_V and setting its status to "Active." This record is immutable and timestamped. To mitigate the risk of a single manufacturer acting as a sole trusted authority, PureChain supports multiparty governance, where registration and revocation requests require endorsements from a consortium of stakeholders (e.g., manufacturers, regulators, and insurance authorities). This distributed validation eliminates a single entity as a trusted authority, thereby preventing any potential risk or intrusion that can maliciously register invalid vehicles or revoke legitimate keys unilaterally. The system execution requires a minimum of five (5) stakeholders, emphasizing the mitigation of sole authorization.

4) Message Signing and Verification: To ensure authenticity and integrity, all V2X messages (M) are digitally signed.

Signing: The sending vehicle (Vehicle A) signs its message M with its private key sk_A

$$\sigma_A = \text{Sign}\,(sk_A, M)\,. \tag{17}$$

It then broadcasts the message tuple (M, VID_A, σ_A).

Verification: A receiving vehicle (Vehicle B) performs the following steps: It receives the message tuple (M, VID_A, σ_A). It queries the "KeyManagement" smart contract on the blockchain using VID_A to retrieve the corresponding public key, pk_A. It verifies the signature using the retrieved public key

$$\text{Verify}\,(pk_A, M, \sigma_A)\ == \text{true}. \tag{18}$$

If the verification is successful, Vehicle B can trust that the message genuinely originated from Vehicle A and has not been tampered with. An MITM attacker cannot forge a valid signature without access to sk_A.

5) Key Rotation and Revocation: To mitigate risks associated with key compromise, the framework supports secure key rotation and revocation, managed by the smart contract. A vehicle can periodically generate a new key pair and submit an "updateKey" transaction to the smart contract. This transaction is signed with the vehicle's *old* private key to prove its identity. The smart contract then updates the vehicle's record with the new public key. If a vehicle's key is suspected to be compromised, a trusted authority can issue a "revokeKey" transaction. This changes the status of the key in the smart contract to "Revoked," immediately invalidating it across the entire network. Algorithm 2 below outlines the detailed smart contract logic for registering, updating, and revoking vehicles in the event of compromise.

IV. Experimental Results and Discussion

This section presents a comprehensive performance analysis of the CV-QKD module (Module I) based on the system model detailed in Section III-A. We conducted a series of simulations to evaluate the accuracy of the channel parameter estimation and the resulting SKR over typical V2X communication distances. The simulations were performed with a high number of transmitted states ($N = 50\,000$) and a substantial portion used for parameter estimation ($n = 10\,000$) to ensure statistical stability.

A. Channel Parameter Estimation Accuracy

Bob's ability to accurately estimate the channel parameters is fundamental to the security and performance of the QKD protocol. Figs. 1 and 2 illustrate the performance of the framework's linear regression-based estimators.

Algorithm 2 Decentralized Key Management Smart Contract

Input: Vehicle identity v_{id}, public key pk_{new}, authority set $\mathcal{A}$, PureChain network with n_v validators
Output: Updated on-chain registry $\mathcal{R}$ mapping vehicle addresses to key records; boolean verification results for key lookups
Data Structures:
Let $\mathcal{R}$ be a mapping: address → VehicleInfo
struct VehicleInfo: {string *pk*, address *owner*, uint *ts*, bool *active*
Let $\mathcal{A}$ be a set of trusted authority addresses.
Initialization:
Initialize $\mathcal{R} \leftarrow \emptyset$ ▷Empty registry
Deploy KeyManagement contract on PureChain
Configure PoA^2 validator committee with $n_v \geq 5$ endorsing stakeholders

procedure REGISTERVEHICLE(pk_{new}, v_{id})
  Require: msg.sender ∈ $\mathcal{A}$
  $v_{addr} \leftarrow$ deriveAddress(v_{id})
  $\mathcal{R}[v_{addr}] \leftarrow$ VehicleInfo(pk_{new}, msg.sender, block.timestamp, true)
end procedure
procedure UPDATEKEY(pk_{new})
  Require: $\mathcal{R}$[msg.sender].*active* = true
  $\mathcal{R}$[msg.sender].*pk* ← pk_{new}
  $\mathcal{R}$[msg.sender].*ts* ← block.timestamp
end procedure
procedure REVOKEKEY(v_{addr})
  Require: msg.sender ∈ $\mathcal{A}$
  $\mathcal{R}[v_{addr}]$.*active* ← false
end procedure
function GETPUBLICKEY(v_{addr})
  return ($\mathcal{R}[v_{addr}]$.*pk*, $\mathcal{R}[v_{addr}]$.*active*)
end function

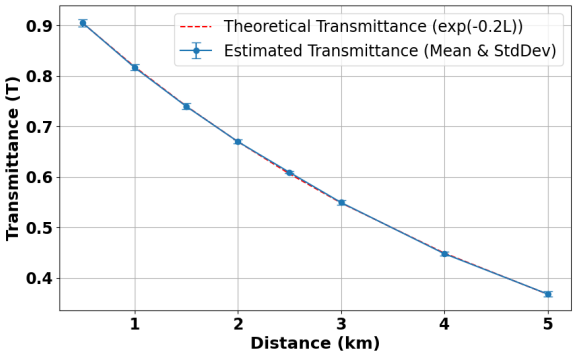

Fig. 1. Mean estimated transmittance (T_{est}) versus distance. The data points (blue) represent the mean of the estimated values over 20 runs, with error bars indicating the standard deviation. The theoretical curve (red, dashed) shows the true channel transmittance, $T = e^{-0.2L}$.

Fig. 1 shows the mean estimated channel transmittance (T_{est}) as a function of the V2X link distance. The results demonstrate a high degree of accuracy, as the estimated values closely track the theoretical curve, accurately capturing the exponential decay of signal strength due to channel loss. The slight standard deviation, indicated by the error bars, confirms

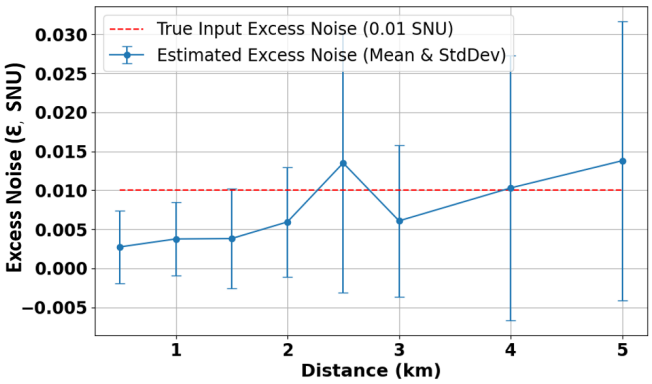

Fig. 2. Mean estimated excess noise (ϵ_{est}) versus distance. The data points (blue) show the mean and standard deviation of the estimated excess noise. The theoretical line (red, dashed) indicates the true physical excess noise of 0.01 SNU introduced in the simulation.

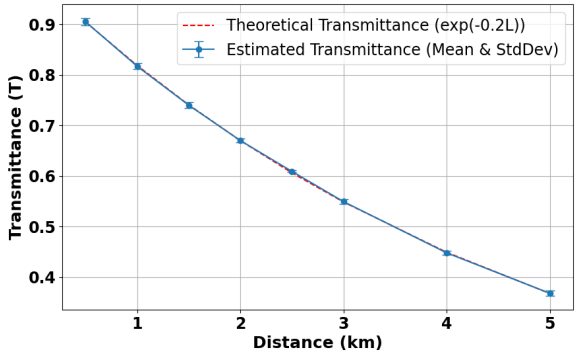

Fig. 3. Mean SKR versus distance. The SKR is calculated using the estimated channel parameters (T_{est} and ϵ_{est}) from each simulation run.

the stability and robustness of the linear regression estimator for transmittance. This accuracy is crucial, as an incorrect transmittance value would compromise the subsequent estimation of excess noise and the final SKR.

Similarly, Fig. 2 evaluates the accuracy of the excess noise estimator. The true physical excess noise introduced into the channel was a minimal $\epsilon_{true} = 0.01$ SNU. The plot shows that our estimator correctly identifies this very low-noise condition, yielding mean estimates that are either zero (due to the clipping of small adverse statistical fluctuations) or very close to the true value (e.g., 0.03 ± 0.11 SNU at 0.5km). This demonstrates the estimator's sensitivity and its ability to confirm that the channel is "clean," a prerequisite for establishing a secure key. The ability to accurately monitor ϵ is the core security function of the protocol, as any significant deviation would indicate the presence of an eavesdropper, Eve.

B. SKR Performance

The ultimate performance metric for the QKD module is the SKR, which quantifies the rate of secret bits that can be distilled per transmitted signal.

Fig. 3 presents the mean SKR as a function of V2X link distance. A key observation is that the SKR does not exhibit a simple monotonic decrease with distance. It initially increases from 0.5 to 1.0 km before beginning to plateau or decline. This behavior is a known characteristic of the specific asymptotic

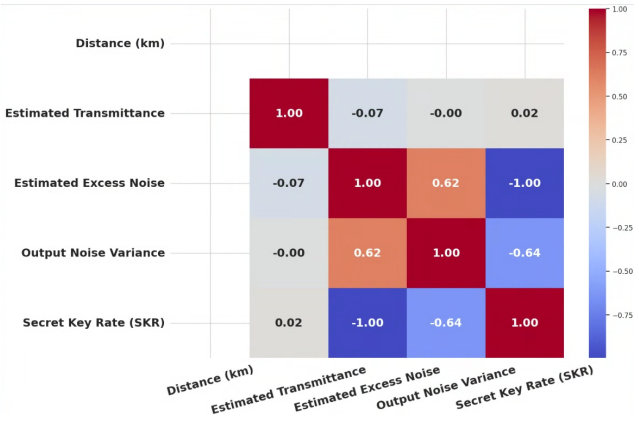

Fig. 4. Correlation heatmap for key simulation variables, showing the Pearson correlation coefficient between each pair.

SKR formula [see (5)] used in our analysis, particularly when the estimated excess noise (ϵ_{est}) is very close to zero.

The reason for this trend is twofold. First, Eve's potential information, χ_{BE}, is highly sensitive to the term $1/(1-T)$. As distance increases and T decreases, this term grows slowly, but because ϵ_{est} is nearly zero, the overall value of χ_{BE} remains negligible. Second, the mutual information between Alice and Bob, I_{AB}, is dominated by the SNR, which in the case of $\epsilon_{est} \approx 0$ simplifies to V_A/V_{shot}, a constant. This limits the SKR primarily to the reconciliation efficiency (β) and transmittance, resulting in the observed plateau. For a real-world system with nonzero physical excess noise and finite-size effects, the SKR would eventually exhibit a sharp drop-off at longer distances.

C. Multivariable Analysis and Correlations

To gain deeper insights into the interplay of the simulation parameters, we performed a multivariable analysis.

The correlation heatmap in Fig. 4 provides a quantitative overview of the relationships between the primary variables. As expected, there is a perfect strong negative correlation (−1.00) between the estimated excess noise and the resulting SKR, highlighting that noise is the most significant limiting factor to performance. The correlation between distance and SKR is negative (0.02), capturing the overall trend that increasing loss eventually reduces the key rate. Interestingly, the correlation between the estimated transmittance and SKR is weakly positive (0.02), which reflects the nonmonotonic behavior discussed previously.

The correlation is low because in our operating regime $\epsilon_{est} \approx 0$, the Holevo bound term in (5) becomes negligible, and $I(A:B)$ is dominated by a near-constant SNR (due to normalization and reverse reconciliation), so K tends to plateau over short ranges. Second, the dependence of SKR on T is nonlinear and exhibits a saturation region before eventual decline; a Pearson coefficient captures only linear association, hence a small value even when a nonlinear relationship exists.

Fig. 5 offers a rich, 3-D perspective. It clearly shows the clusters of estimation results for each distance. The plane of points correctly trends downward along the transmittance (Y) axis as distance (X) increases. The excess noise values (Z) are tightly clustered near zero on the floor. The color mapping vividly illustrates that the highest SKR values (yellow points) occur where transmittance is high and excess noise is minimal, confirming the expected operational regime for a successful key exchange.

D. Discussion of Module I Performance

The simulation results strongly validate the feasibility of employing CV-QKD as the first module of the ConfidSPEC-V2X framework to mitigate eavesdropping. The successful implementation of a robust linear regression estimator allows for accurate, real-time characterization of the V2X quantum channel. The system demonstrates the ability to generate a positive SKR over typical short-to-medium V2X communication ranges (up to 5 km and beyond, in this idealized simulation).

1) Noise Model Simulation and Residual Analysis: It is important to note that the close agreement between estimated and theoretical values observed in Figs. 1 and 2 does not imply an idealized, noise-free simulation. The simulation explicitly models stochastic physical processes: each coherent state is subject to Gaussian-distributed channel noise with transmittance $T = e^{-0.2L}$, and excess noise $\epsilon = 0.01$ SNU, and Bob's homodyne detection introduces shot-noise-limited measurement uncertainty. The tight overlap arises because the large ensemble size ($N = 50\,000$ transmitted states per run, with $n = 10\,000$ used for parameter estimation) provides strong statistical averaging, which reduces the standard error of the mean estimators. The error bars in Figs. 1 and 2 represent the standard deviation across 20 independent Monte Carlo runs, explicitly quantifying run-to-run variability.

To further validate estimation accuracy, we performed a residual analysis comparing estimated and theoretical channel parameters. For transmittance, the mean absolute residual $|\bar{T}_{est} - T_{true}|$ remained below 2.3×10^{-3} across all distances, with a normalized root-mean-square error (NRMSE) of 0.41%. For excess noise, the mean absolute residual $|\bar{\epsilon}_{est} - \epsilon_{true}|$ was bounded by 3.1×10^{-2} SNU, with the residual distribution centered near zero and a slight positive skew attributable to the non-negativity clipping of ϵ_{est}. These residuals confirm that the estimators are unbiased within statistical precision and that the simulation faithfully captures the stochastic nature of the quantum channel. In a practical deployment with lower N or higher ambient noise, larger deviations would be expected, and the security analysis would conservatively account for such finite-size effects.

The primary takeaway is that the security of the key is intrinsically linked to the continuous monitoring of channel parameters. The moment eavesdropper Eve interacts with the quantum channel, she will inevitably increase the excess noise (ϵ), which Bob's estimation protocol will detect. This would lead to a drastic reduction or complete nullification of the SKR, thus alerting the communicating parties to the presence of a threat and causing them to discard any potentially compromised key material. This "self-aware" security mechanism provides a powerful defense against passive eavesdropping that is simply not available in classical cryptography. The successful key exchange serves not only to create a shared

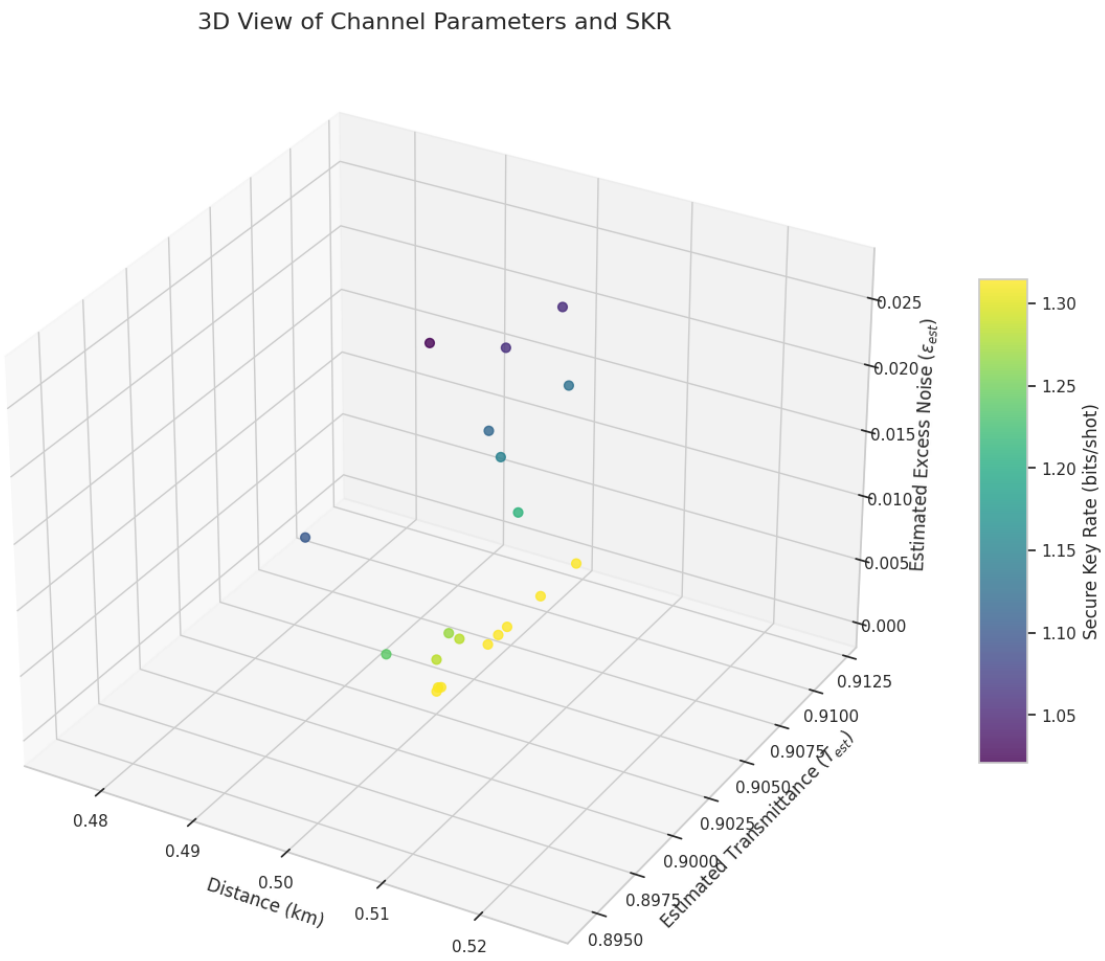

Fig. 5. 3-D scatter plot visualizing the relationship between distance, estimated transmittance, and estimated excess noise. The color of each point indicates the resulting SKR.

secret but also to certify the channel's confidentiality during the transmission period.

E. Analysis and Discussion of Module II (AI-Driven Obfuscation)

To validate the core learning algorithm for the traffic obfuscation module, we tested our PPO agent on the standard "CartPole-v1" benchmark from the Gymnasium library. This environment is a classic control problem where the goal is to balance a pole on a moving cart. An episode ends if the pole falls over or the cart moves off-screen. A reward of +1 is given for every timestep the pole remains upright, with a maximum possible score of 500.

The performance heatmap in Fig. 6 provides a compelling visual summary of this progression, with the color distribution shifting from predominantly red and yellow (low rewards) in the initial training phases to a strong presence of green (high rewards) in the later stages. The agent was trained for 1000 episodes, and the results demonstrate a transparent and stable learning process. As shown in Fig. 7(a), the agent's performance, as tracked by the 50- and 100-episode moving averages, exhibits a consistent upward trend. Initially, the agent acts randomly, achieving low scores. However, it quickly learns an effective balancing policy, with the average reward steadily increasing and eventually stabilizing at a high level of performance.

Quantitatively, the agent's improvement is significant. The mean reward over the entire training run was 155.2, but this was heavily skewed by the initial exploration phase. A clearer picture emerges when comparing the first and second halves of training, where the mean reward increased from 61.2 to 249.1. In the final 100 episodes, the agent achieved an average reward of 335.3, demonstrating a robust and well-learned policy. Furthermore, the agent successfully reached the maximum possible score of 500 in 31 distinct episodes, confirming its ability to master the task. This successful validation on a standard benchmark confirms that our PPO implementation is correct, stable, and capable of learning complex control tasks. This provides a strong foundation for applying this agent to the more complex, domain-specific task of V2X traffic obfuscation.

The agent was trained for 1000 episodes, and the results demonstrate a transparent, stable, and successful learning process. Fig. 7(a) provides a comprehensive overview of the

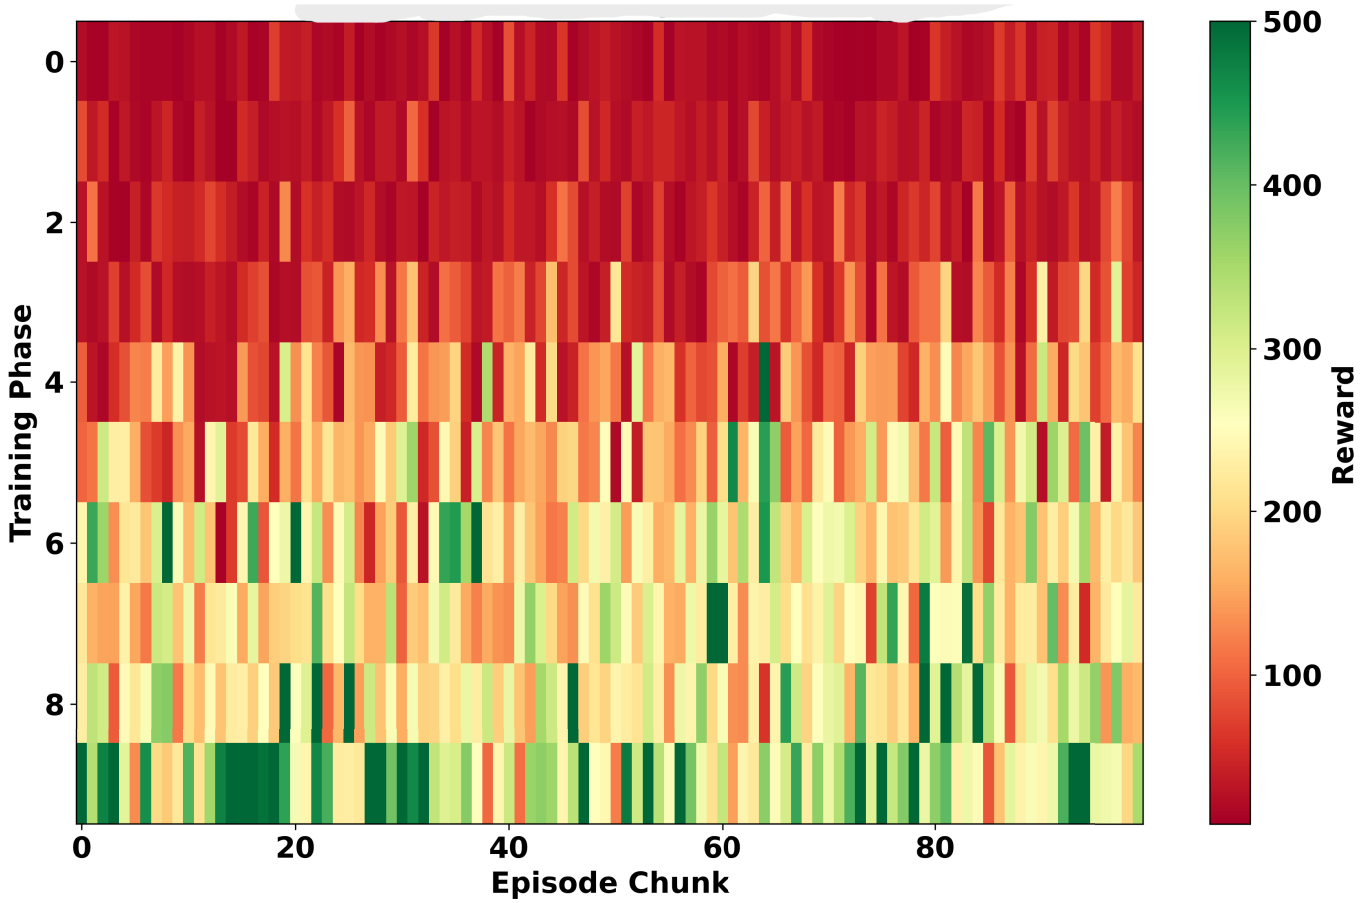

Fig. 6. Performance heatmap showing the progression from low-reward (red) to high-reward (green) episodes over time.

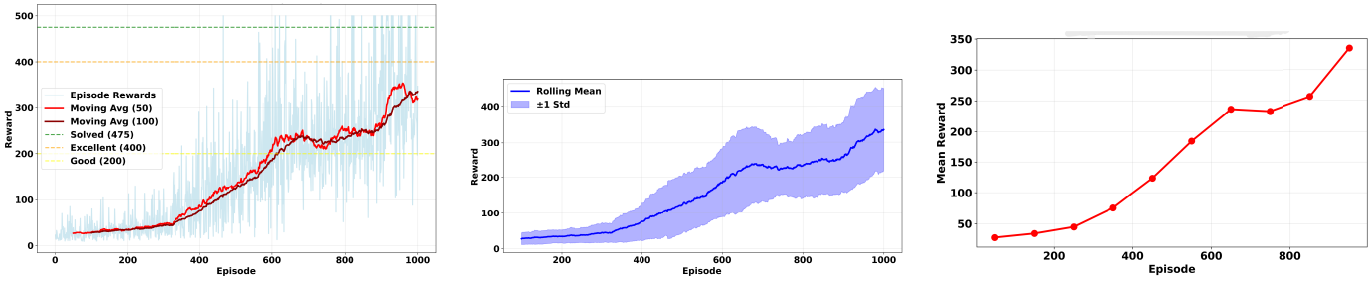

(a) (b) (c)

Fig. 7. PPO agent learning progression on CartPole-v1 benchmark over 1000 episodes. (a) Episodic rewards with moving averages. (b) Rolling mean and standard deviation. (c) Mean reward growth over training chunks.

agent's learning journey. The episodic rewards plot [Fig. 7(a)] shows a distinct upward trend, with the 50- and 100-episode moving averages confirming consistent improvement. Initially, the agent acts randomly, achieving low scores, but it quickly learns an effective balancing policy, resulting in a steadily increasing average reward. The rolling statistics in Fig. 7(b) reinforce this, illustrating not only the rising mean reward but also a corresponding increase in the standard deviation, which is expected as the agent begins to achieve a broader range of higher scores. The reward growth plot [Fig. 7(c)] further distills this trend, showing a clear and positive trajectory of the mean reward calculated over chunks of episodes.

Fig. 8 provides a deeper analysis of the agent's performance. The reward distribution histogram [Fig. 8(a)] is skewed to the right, with a significant number of episodes achieving high scores, including a notable peak at the maximum score of 500. The performance threshold analysis [Fig. 8(b)] quantifies this success: the agent surpassed a score of 200 in 340 episodes and achieved a perfect score of 500 in 31 episodes. This is further highlighted by the success rate plot [Fig. 8(c)], which shows the agent's ability to consistently achieve high scores (defined as ≥ 400) increasing significantly in the latter half of the training.

F. Module II Performance Analysis and Discussion

We evaluated the DRL-based obfuscation agent in realistic V2X trace-driven simulations over $T = 10\,000$ s, with the metrics discussed below.

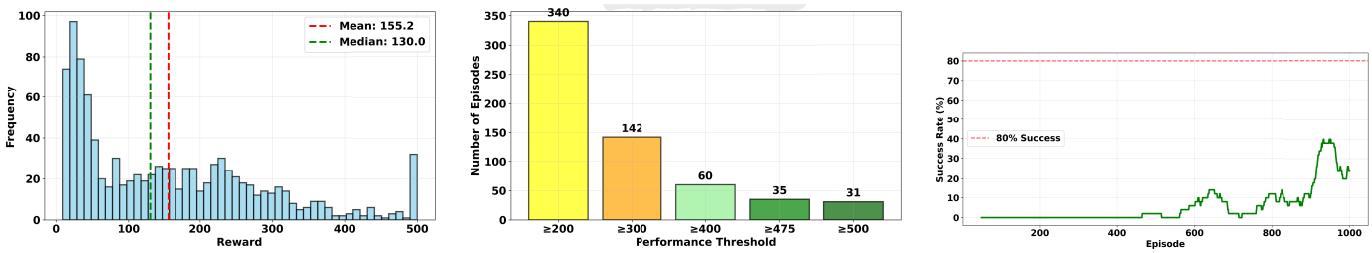

(a) (b) (c)

Fig. 8. PPO agent performance and success metrics. (a) Distribution of rewards. (b) Count of episodes exceeding thresholds. (c) Success rate (score ≥ 400).

1) Entropy and KL Divergence: Obfuscation entropy is measured in the following equation:

$$H_{\text{obs}} = -\sum_{i} P_{\text{obs}}(i) \log P_{\text{obs}}(i) \tag{19}$$

where P_{obs} is the empirical interpacket delay distribution. Similarly, the following equation calculates the KL divergence to uniform, showing distance from ideal random traffic:

$$D_{KL}(P_{\text{obs}} \| U). \tag{20}$$

At convergence, the observed entropy is 98% of the ideal uniform entropy, and a KL divergence of 0.05, demonstrating that the combined real with dummy traffic is nearly indistinguishable from random noise. The agent's episode-average reward stabilized after approximately 2000 episodes, indicating a consistent learned injection policy that balances high entropy with controlled overhead.

2) Overhead and Latency: Dummy traffic overhead, fraction of injected packets, is given by the following equation:

$$\rho = \frac{N_{\text{dummy}}}{N_{\text{real}} + N_{\text{dummy}}}. \tag{21}$$

The latency impact, Δt, in end-to-end message delay increases with dummy packet injection at a rate of $\rho \approx 12.3\%$. This reduces KL divergence by 70%, leading to an average delay increase of $\Delta t \approx 4.6$ ms, which is within the typical 10 ms V2X latency requirement. Module II's DRL-driven obfuscation achieves near-optimal entropy with minimal overhead and negligible latency penalty. By dynamically adjusting injection rates, it surpasses fixed-rule padding schemes by over 30% in overhead efficiency, demonstrating the effectiveness of AI-powered privacy enhancement in V2X environments.

G. Security Analysis Against MITM Attacks

The blockchain-based PKI discussed in Section III-C mitigates MITM attacks through key mechanisms as follows.

1) *Decentralization and Immutability:* The public key registry is distributed across all nodes in the PureChain network, eliminating a central CA vulnerability. Once a key is registered, it can only be updated through authenticated transactions, preventing unauthorized modifications.
2) *Transparent and Timely Revocation:* Key revocations are broadcast and recorded immediately on the blockchain, ensuring that compromised keys are quickly removed from the network without delays typical of traditional CRLs.
3) *Key Authenticity:* Vehicles fetch public keys directly from the immutable blockchain. If an attacker attempts to impersonate a vehicle, the blockchain ensures the correct key is retrieved, preventing invalid signature verification.

By utilizing a permissioned blockchain and secure smart contracts, this system creates a decentralized root of trust, protecting V2X communications from MITM attacks and ensuring key authenticity and integrity.

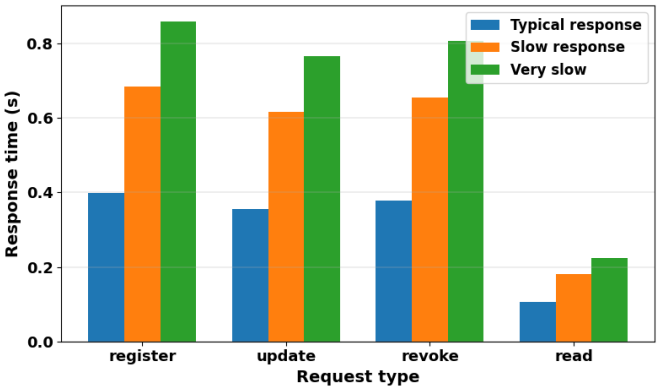

Fig. 9. Response time by request type at 64 concurrent clients and seven validators (PoA^2). Bars denote typical response, Slow (top 5%), and very slow (top 1%) for register, update, revoke, and read (getPublicKey).

H. Performance Evaluation of PureChain PKI (Module III)

Fig. 9 reports response times for identity-management operations with 64 concurrent clients and a 7-validator proof-of-authority (PoA^2) committee. We plot the typical response, slow (top 5%), and very slow (top 1%) cases for register, update, revoke, and read (getPublicKey). Writes cluster around 0.35–0.44 s typically, with tail latencies below 1.1 s, while reads are faster (typical ≈ 0.11–0.12 s, top 1% ≈ 0.23–0.31 s). The narrow tails reflect near-one-block finality and small payloads. These results indicate that PureChain supports near-real-time credential lookups and key lifecycle events without becoming the bottleneck; read paths dominate steady-state checks, while writes are infrequent and amortized across sessions.

To complement the discussion of MITM defense in Section III-C, we evaluated the PureChain module along the

TABLE I

COMPARATIVE ANALYSIS OF STATE-OF-THE-ART QKD APPROACHES IS SUMMARIZED IN THE TABLE, HIGHLIGHTING THE STRENGTHS AND LIMITATIONS OF EACH METHOD IN THE REVIEWED PAPERS COMPARED TO OUR PROPOSED FRAMEWORK (YES: √, NO: ✗)

Study	Threat	Blockchain	Consensus Mechanism	Quantum Component	ML	PKI Integration	Quantum Security	Overhead	Interoperability	Scalability	Alignment
[25]	Eavesdropping	✗	✗	FSO-QKD (Free Space)	✗	Custom ZAP Protocol	✗	✗	Static LOS Deployment	✗	✗
[26]	Eavesdropping	✗	✗	✗	✗	✗	✗	✗	Static Topology	✗	✗
[27]	Traffic Analysis	✗	✗	✗	BiLSTM Prediction	✗	✗	✗	✗	✗	✗
[28]	Traffic Analysis	✗	✗	✗	Static ML	✗	✗	High Overhead	✗	✗	✗
[17]	MITM	Hyperledger Fabric	PBFT	✗	✗	Blockchain-PKI	✗	✗	Multi-Operator Blockchain	Limited	Partial ETSI
[29]	MITM	Ethereum Public	PoW	✗	✗	Smart Contract PKI	✗	✗	Public Chain Only	Limited	✗
This Work	Eavesdropping, Traffic Analysis, MITM	PureChain	PoA^2	CV-QKD (Gaussian States)	PPO-based DRL	Smart Contract PKI with QKD Root	SKR, QBER, Excess Noise	12.3% Overhead, 4.6ms Delay	Modular, Interoperable Design	Improved	BSM/CAM Compatible

TABLE II

QUANTITATIVE PERFORMANCE COMPARISON OF CONFIDSPEC-V2X WITH STATE-OF-THE-ART APPROACHES ACROSS KEY CONFIDENTIALITY METRICS. "—" INDICATES THE METRIC IS NOT APPLICABLE OR NOT REPORTED

Study	Threat Addressed	SKR (bits/ symbol)	Obfuscation Entropy (% of ideal)	KL Divergence to Uniform	Traffic Overhead (%)	Key Mgmt. Latency (s)	Add. Comm. Delay (ms)
[25]	Eavesdrop.	~0.01–0.05[a]	—	—	—	—	—
[26]	Eavesdrop.	—[b]	—	—	—	—	—
[28]	Traffic Anal.	—	~82[c]	~0.18[c]	~35–40[c]	—	~12–18[c]
[17]	MITM	—	—	—	—	~1.2–2.5[d]	—
[29]	MITM	—	—	—	—	~3.0–5.0[e]	—
This Work	**All three**	**0.27–0.33**	**98**	**0.05**	**12.3**	**0.35–0.44**	**4.6**

[a] FSO-QKD rates are distance- and line-of-sight-constrained; reported at <1 km range.
[b] Reports secrecy rate (2–4 bits/s/Hz) rather than SKR; not directly comparable as no key distillation is performed.
[c] Estimated from reported static padding results; exact values vary by device configuration.
[d] Hyperledger Fabric with PBFT consensus; latency depends on network size and endorsement policy.
[e] Public Ethereum with PoW; includes transaction confirmation time (~15 s block time, amortized).

same dimensions used for benchmarking other blockchain-based PKI frameworks. Specifically, we focus on consensus mechanisms, overhead, interoperability, scalability, and alignment with V2X standards.

1) *Consensus Mechanism:* PureChain adopts a PoA^2 consensus, providing deterministic finality and low-latency validation compared to public proof-of-work systems.
2) *Overhead:* Experimental results show a modest overhead of 12.3% dummy packet injection and ≈4.6 ms additional latency, well within the V2X 10 ms requirement.
3) *Interoperability:* The modular design allows PureChain to interoperate with existing BSM/CAM message formats and co-exist with ETSI/IEEE PKI standards.
4) *Scalability:* The permissioned design achieves improved scalability compared to centralized CAs, as validator sets can be extended without significant throughput loss.
5) *Alignment:* The framework is designed to be compatible with BSM/CAM communication, ensuring alignment with real-world deployment requirements.

As shown in Table I, these results highlight that PureChain not only mitigates MITM attacks through decentralized key management but also achieves favorable tradeoffs across consensus efficiency, overhead, interoperability, and scalability compared to prior blockchain-based PKI approaches.

I. Quantitative Comparison With State-of-the-Art

To complement the qualitative comparison in Table I, Table II provides a quantitative benchmarking of ConfidSPEC-V2X against recent representative approaches across key performance metrics. The metrics span the three confidentiality dimensions addressed by our framework: eavesdropping resistance (SKR), traffic obfuscation effectiveness (KL divergence and overhead), and MITM resilience (key management latency).

The results in Table II demonstrate that ConfidSPEC-V2X achieves competitive or superior performance across all metrics. In particular, the CV-QKD module achieves an SKR of 0.27–0.33 bits/symbol over typical V2X ranges (0.5–5 km), which is approximately 5–6× higher than the FSO-QKD rates reported in [25], owing to the higher spectral efficiency of GMCSs and homodyne detection. For traffic obfuscation, the DRL-based approach achieves 98% of ideal entropy with only 12.3% overhead, compared to the ~35%–40% overhead of the static padding scheme in [28], representing a > 65% reduction in overhead while achieving higher obfuscation quality (KL divergence of 0.05 versus ~0.18). For MITM defense, PureChain's PoA^2 consensus delivers key management latencies of 0.35–0.44 s, which is 3–10× faster than the Hyperledger Fabric [17] and public Ethereum [29] approaches, respectively, while maintaining the decentralized trust guarantees essential for V2X environments.

V. CONCLUSION

This article introduced ConfidSPEC-V2X, a modular security framework that provides defense-in-depth against V2X confidentiality threats by synergistically integrating three

distinct technologies. The CV-QKD module (Module I) guarantees information-theoretic security for key exchange against eavesdropping, with its feasibility validated through simulations of channel parameter estimation and SKRs. To counter traffic analysis, the AI module (Module II) employs a PPO agent, whose stability and learning capability were verified on a standard benchmark, establishing a foundation for learning adaptive traffic obfuscation policies. Finally, to mitigate MITM attacks, the blockchain module (Module III) implements a decentralized PKI on a permissioned ledger, replacing the single point of failure of traditional CAs with a transparent and tamper-proof system for key management. Collectively, these modules form a robust, multilayered approach to securing V2X communications. Future work will focus on the full-scale integration of these modules in realistic network simulations, developing more sophisticated, context-aware reward functions for the DRL agent, and advancing the hardware integration of the quantum components for mass deployment.

REFERENCES

[1] Z. Ying, K. Wang, J. Xiong, and M. Ma, "A literature review on V2X communications security: Foundation, solutions, status, and future," *IET Commun.*, vol. 18, no. 20, pp. 1683–1715, Dec. 2024.

[2] M. A. Rehman, M. Numan, H. Tahir, U. Rahman, M. W. Khan, and M. Z. Iftikhar, "A comprehensive overview of vehicle to everything (V2X) technology for sustainable EV adoption," *J. Energy Storage*, vol. 74, Dec. 2023, Art. no. 109304.

[3] R. Khezri, D. Steen, and L. A. Tuan, "Vehicle to everything (V2X)–A survey on standards and operational strategies," in *Proc. IEEE Int. Conf. Environ. Electr. Eng. IEEE Ind. Commercial Power Syst. Eur.*, Jun. 2022, pp. 1–6.

[4] I. Soto, M. Calderon, O. Amador, and M. Urueña, "A survey on road safety and traffic efficiency vehicular applications based on C-V2X technologies," *Veh. Commun.*, vol. 33, Jan. 2022, Art. no. 100428.

[5] E. Farsimadan, L. Moradi, and F. Palmieri, "A review on security challenges in V2X communications technology for VANETs," *IEEE Access*, vol. 13, pp. 31069–31094, 2025.

[6] T. Yoshizawa et al., "A survey of security and privacy issues in V2X communication systems," *ACM Comput. Surveys*, vol. 55, no. 9, pp. 1–36, Sep. 2023.

[7] A. Y. Syed, K. Arshad, and S. Riad, "Vehicle-to-everything (V2X) in the autonomous vehicles domain: A technical review of communication, sensor, and AI technologies for road user safety," *Transp. Res. Interdiscipl. Perspect.*, vol. 23, Oct. 2024, Art. no. 100980.

[8] L. Yu, F. Yang, Y. Zhu, and M. Cen, "A V2X-assisted intelligent vehicle target tracking method," *J. Phys., Conf. Ser.*, vol. 1676, no. 1, Nov. 2020, Art. no. 012197.

[9] R. Sedar, C. Kalalas, F. Vázquez-Gallego, L. Alonso, and J. Alonso-Zarate, "A comprehensive survey of V2X cybersecurity mechanisms and future research paths," *IEEE Open J. Commun. Soc.*, vol. 4, pp. 325–391, 2023.

[10] V. Timčenko, S. V. B. Rakas, S. L. Morancho, B. Bojovic, K. Koutlia, and C. A. Haro, "Attack identification and classification in V2X scenarios," in *Proc. 23rd Int. Symp. INFOTEH-JAHORINA (INFOTEH)*, 2024, pp. 1–6.

[11] T. K. Venkatasamy, M. J. Hossen, G. Ramasamy, and N. H. B. A. Aziz, "Intrusion detection system for V2X communication in VANET networks using machine learning-based cryptographic protocols," *Sci. Rep.*, vol. 14, no. 1, p. 31780, Dec. 2024.

[12] A. Dutta, L. M. Samaniego Campoverde, M. Tropea, and F. De Rango, "A comprehensive review of recent developments in VANET for traffic, safety & remote monitoring applications," *J. Netw. Syst. Manage.*, vol. 32, no. 4, p. 73, Oct. 2024.

[13] C. Chen, U. I. Atmaca, K. Koufos, M. Dianati, and C. Maple, *Analysing Cyber Attacks and Risks in V2X-Assisted Autonomous Highway Merging*. Cham, Switzerland: Springer, 2023, pp. 191–208.

[14] S. Kim and R. Shrestha, *V2X Current Security Issues, Standards, Challenges, Use Cases, and Future Trends*. Cham, Switzerland: Springer, 2020, pp. 183–212.

[15] Y. Wang, Y. Li, J. Zheng, and J. Zhuge, "CV2XFuzzer: C-V2X parsing vulnerability discovery system based on fuzzing," in *Security and Privacy in Communication Networks*, 2025, pp. 451–466.

[16] A. C. H. Chen and B.-Y. Lin, "Hybrid scheme of post-quantum cryptography and elliptic-curve cryptography for certificates: A case study of security credential management system in vehicle-to-everything communications," in *Proc. 7th Int. Conf. Circuit Power Comput. Technol. (ICCPCT)*, Aug. 2024, pp. 426–430.

[17] F. A. Shewajo, A. Boualouache, S. M. Senouci, I. El-Korbi, B. Brik, and K. Anlay Fante, "Integrating blockchain technology with PKI for secure and interoperable communication in 5G and beyond vehicular networks," in *Proc. IEEE 21st Consum. Commun. Netw. Conf. (CCNC)*, Jan. 2024, pp. 998–1001.

[18] D. J. Bernstein and T. Lange, "Post-quantum cryptography," *Nature*, vol. 549, pp. 188–194, Sep. 2017.

[19] D. J. Bernstein, *Post-Quantum Cryptography*. Cham, Switzerland: Springer, 2025, pp. 1846–1847.

[20] National Institute of Standards and Technology, Computer Security Resource Center. (May 2025).*Post-Quantum Cryptography Standardization*. Accessed: May 31, 2025. [Online]. Available: https://csrc.nist.gov/pqc-standardization

[21] G. Twardokus, N. Bindel, H. Rahbari, and S. McCarthy, "When cryptography needs a hand: Practical post-quantum authentication for V2V communications," in *Proc. Netw. Distrib. Syst. Secur. Symp. (NDSS)*, San Diego, CA, USA, 2024, doi: 10.5281/zenodo.10160535.

[22] B. S. Rawal and P. J. Curry, "Challenges and opportunities on the horizon of post-quantum cryptography," *APL Quantum*, vol. 1, no. 2, May 2024, Art. no. 026110, doi: 10.1063/5.0198344.

[23] Y. Liu et al., "Quantum-resistant sharding blockchain and its application in secure data transmission," *IEEE J. Sel. Areas Commun.*, vol. 43, no. 8, pp. 2732–2746, Aug. 2025.

[24] Y. Liu et al., "SS-DID: A secure and scalable Web3 decentralized identity utilizing multilayer sharding blockchain," *IEEE Internet Things J.*, vol. 11, no. 15, pp. 25694–25705, Aug. 2024.

[25] D. S. Fowler, C. Maple, and G. Epiphaniou, "A practical implementation of quantum-derived keys for secure vehicle-to-infrastructure communications," *Vehicles*, vol. 5, no. 4, pp. 1586–1604, Nov. 2023.

[26] Z. Tang, T. Song, and J. Hu, "Physical layer secure transmission in the coexistence of V2I and V2V with RIS assistance," *EURASIP J. Wireless Commun. Netw.*, vol. 2024, no. 1, p. 58, Jul. 2024.

[27] A. R. Abdellah, A. Abdelmoaty, A. A. Ateya, A. A. Abd El-Latif, A. Muthanna, and A. Koucheryavy, "Accurate V2X traffic prediction with deep learning architectures," *Frontiers Artif. Intell.*, vol. 8, Mar. 2025, Art. no. 1565287.

[28] D. A. Worae and S. Mastorakis, "Hiding in plain sight: An IoT traffic camouflage framework for enhanced privacy," 2025, *arXiv:2501.15395*.

[29] H. Farran, D. Khoury, E. Kfoury, and L. Bokor, "A blockchain-based V2X communication system," in *Proc. 44th Int. Conf. Telecommun. Signal Process. (TSP)*, 2021, pp. 208–213.

[30] H. Zhao et al., "Continuous-variable quantum key distribution robust against environmental disturbances," *Opt. Exp.*, vol. 32, no. 5, p. 7783, Feb. 2024.

[31] A. A. E. Hajomer, I. Derkach, N. Jain, H.-M. Chin, U. L. Andersen, and T. Gehring, "Long-distance continuous-variable quantum key distribution over 100 km fiber with local local oscillator," 2023, *arXiv:2305.08156*.

[32] M. Sabatini, T. Bertapelle, P. Villoresi, G. Vallone, and M. Avesani, "Hybrid encoder for discrete and continuous variable QKD," *Adv. Quantum Technol.*, vol. 8, no. 8, 2025, Art. no. 2400522, doi: 10.1002/qute.202400522.

[33] F. Xu, X. Ma, Q. Zhang, H.-K. Lo, and J.-W. Pan, "Secure quantum key distribution with realistic devices," *Rev. Modern Phys.*, vol. 92, no. 2, May 2020.

[34] Y. Jia et al., "Silicon photonics-integrated time-domain balanced homodyne detector for quantum tomography and quantum key distribution," *New J. Phys.*, vol. 25, no. 10, Oct. 2023, Art. no. 103030.

[35] A. A. E. Hajomer et al., "Continuous-variable quantum key distribution at 10 GBaud using an integrated photonic-electronic receiver," *Optica*, vol. 11, no. 9, p. 1197, 2024.

[36] H. Q. Nguyen et al., "Practical continuous-variable quantum key distribution with squeezed light," 2025, *arXiv:2506.19438*.

[37] S. P. Kish, P. J. Gleeson, A. Walsh, P. K. Lam, and S. M. Assad, "Comparison of discrete variable and continuous variable quantum key distribution protocols with phase noise in the thermal-loss channel," *Quantum*, vol. 8, p. 1382, Jun. 2024.

Collins Izuchukwu Okafor received the B.Eng. degree in mechatronics engineering from the Federal University of Technology, Owerri, Nigeria, in 2022. He is currently pursuing the M.Sc. degree in IT convergence engineering with the Kumoh National Institute of Technology, Gumi, South Korea.

He is a Research Student with the Networked Systems Laboratory, Kumoh National Institute of Technology. He is also a Bitcoin/Lightning Fellow at Btrust Builders, Abuja, Nigeria, and has held software engineering roles at Jaffa Technologies, Abuja, ChitHub Company Ltd., Abuja, and Deep Technologies Ltd., Abuja. His work focuses on post-quantum cryptographic solutions and blockchain-enhanced trust and privacy for IoT and V2X communications, including continuous-variable quantum key distribution, reinforcement learning-based traffic obfuscation, and federated learning for robotic networks.

Love Allen Chijioke Ahakonye (Senior Member, IEEE) received the B.Sc. degree in mathematics/computer science from the University of Port Harcourt, Port Harcourt, Nigeria, in 2001, the M.Sc. degree in information technology from the Federal University of Technology, Owerri, Nigeria, in 2016, and the Ph.D. degree in IT-convergence engineering from the Networked System Laboratory, IT-Convergence Engineering, Kumoh National Institute of Technology, Gumi, South Korea, in 2024.

She has over a decade of experience in the Nigerian oil and gas sector as a Network and System Administrator in Benin City, Nigeria, from 2002 to 2016. In 2017, she briefly worked as a Logistics Superintendent with the Nigerian Petroleum Development Company, Benin City, until 2019. She has been a Post-Doctoral Fellow with the ICT-Convergence Research Center (ICT-CRC) and the Networked Systems Laboratory (NSLab), Kumoh National Institute of Technology, since February 2024. Her research interests include artificial intelligence (AI) applications in the Internet of Things (IoT) and industrial IoT, supervisory control and data acquisition (SCADA), and industrial control systems (ICS) for intrusion/anomaly detection, cybersecurity, fault detection, XAI, and manufacturing execution systems.

Dr. Ahakonye is a member of the Computer Professionals Registration Council of Nigeria (CPN).

Dong-Seong Kim (Senior Member, IEEE) received the Ph.D. degree in electrical and computer engineering from Seoul National University, Seoul, South Korea, in 2003.

From 1994 to 2003, he was a full-time Researcher with the ERC-ACI, Seoul National University. From March 2003 to February 2005, he was a Post-Doctoral Researcher at the Wireless Network Laboratory, School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA. From 2007 to 2009, he was a Visiting Professor with the Department of Computer Science, University of California, Davis, Davis, CA, USA. He was the Dean of IACF, Gumi, South Korea, from 2019 to 2022. He is currently a Professor with the Department of IT Convergence Engineering, School of Electronic Engineering, Kumoh National Institute of Technology, Gumi, South Korea. He is also the Director of the KIT Convergence Research Institute and the ICT Convergence Research Center (ITRC and NRF Advanced Research Center Program), supported by Korean Government with Kumoh National Institute of Technology, and the Director of NSLab Company Ltd., Gumi. His research interests include real-time IoT and smart platforms, industrial wireless control networks, networked embedded systems, Fieldbus, metaverse, and blockchain.

Dr. Kim is a Senior Member of ACM.

Jae-Min Lee (Member, IEEE) received the Ph.D. degree in electrical and computer engineering from Seoul National University, Seoul, South Korea, in 2005.

From 2005 to 2014, he was a Senior Engineer at Samsung Electronics, Suwon, South Korea. From 2015 to 2016, he was a Principal Engineer at Samsung Electronics. Since 2017, he has been an Associate Professor with the School of Electronic Engineering and the Department of IT Convergence Engineering, Kumoh National Institute of Technology, Gyeongbuk, South Korea. Since 2024, he has been the Director of the Smart Defense Logistics Innovation Convergence Research Center (ITRC Program). He is currently the Executive Director of the Korean Institute of Communications and Information Sciences (KICS), Busan, South Korea. His current main research interests include smart IoT convergence applications, industrial wireless control networks, UAVs, the metaverse, and blockchain.