Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
Pending assignment.	n/a	n/a

4. Title and Subtitle	5. Report Date
Enhancing Security and Privacy in Vehicular Networks	February 24, 2026
	6. Performing Organization Code n/a

7. Author(s)	8. Performing Organization Report No.
Sarah Al-Shareeda, Ph.D. https://orcid.org/0000-0001-6337-0453 Fusun Ozguner, Ph.D. https://orcid.org/0009-0001-7516-7089	n/a

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
The Ohio State University Center for Automotive Research 930 Kinnear Road Columbus, OH 43212	n/a
	11. Contract or Grant No. 69A3552348327

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
CARMEN+ University Transportation Center The Ohio State University 930 Kinnear Road Columbus, OH 43212	Final (Aug '23 to Dec '25)
	14. Sponsoring Agency Code USDOT

15. Supplementary Notes

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by a grant from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

16. Abstract

Vehicle-to-Everything (V2X) communication underpins Highly Automated Transportation Systems (HATS), yet scalability is constrained by three coupled challenges: (i) authentication overhead in latency-sensitive communication, (ii) mobility-aware computation offloading under dynamic wireless and infrastructure conditions, and (iii) safety-critical decision-making under behavioral uncertainty. Existing architectures often treat security, offloading, synchronization, and control as loosely coupled components, leading to systematic latency underestimation and degraded responsiveness in dense Cellular-V2X (C-V2X) environments.

This work advances an integrated cyber-layer intelligence framework evaluated through three coordinated tracks that explicitly couple authentication, mobility, synchronization, and learning within a unified decision loop.

Track 1 embeds Group Signatures (GS) and Identity-Based Cryptography (IBC) directly into end-to-end latency modeling and optimization. Results show that authentication introduces a persistent baseline delay and that C-V2X handovers amplify end-to-end latency by up to 47% while nearly doubling authentication frequency in dense regimes. By incorporating these costs into mobility-aware optimization, the framework prevents delay underestimation and stabilizes offloading decisions under scale and handover.

Track 2 extends Digital Twins (DTs) from passive mirrors to predictive coordination surfaces. Prediction-enabled synchronization reduces cyber-physical misalignment by approximately 88-94% and translates into measurable deployment-level gains, with DT-assisted cloud, edge, and hybrid configurations consistently reducing latency relative to reactive baselines across density, task-size, and transmission-rate scaling.

Track 3 addresses behavioral uncertainty using an uncertainty-aware Reinforcement Learning (RL) controller for autonomous safety. The proposed optimization policy achieves a collision rate of 0.0584 at a 10% hesitation probability and maintains bounded degradation as uncertainty increases, demonstrating adaptive safety without reliance on explicit trajectory prediction.

Overall, the results show that secure vehicular intelligence is inherently coupled: authentication and mobility define the latency structure, synchronization fidelity determines decision timeliness, and adaptive control governs safety under uncertainty. By coordinating these dimensions within a predictive cyber layer, the framework improves scalability, robustness, and safety in dense and dynamic vehicular environments, directly supporting USDOT and CARMEN+ priorities for secure and resilient intelligent transportation systems.

17. Keywords		**18. Distribution Statement**	
Connected and Automated Vehicles (CAV), Vehicle-to-Everything (V2X) Communication, Cellular Vehicle-to-Everything (C-V2X), Intelligent Transportation Systems (ITS), Transportation Cybersecurity, Authentication, Edge Computing, Task Offloading, Digital Twins, Reinforcement Learning, Autonomous Vehicle Safety, Uncertainty Quantification		No restrictions.	
19. Security Classif.(of this report) Unclassified	**20. Security Classif.(of this page)** Unclassified	**21. No. of Pages** 56	**22. Price** n/a

Form DOT F 1700.7 (8-72) **Reproduction of completed page authorized**

CARMEN+

Center for Automated Vehicles Research with Multimodal Assured Navigation

USDOT University Transportation Centers Program

Final Report: Enhancing Security and Privacy in Vehicular Networks

P.I.	Project Info:
Fusun Ozguner, Ph.D. https://orcid.org/0009-0001-7516-7089 Sarah Al-Shareeda, Ph.D. https://orcid.org/0000-0001-6337-0453 The Ohio State University Center for Automotive Research	Grant No. 69A3552348327 DUNS: 832127323 EIN #: 31-6025986 Project Effective: Aug 1, 2023 Project End: Dec 31, 2025 Submission: February 26, 2026

Consortium Members:

The Ohio State University College of Engineering | TEXAS The University of Texas at Austin | UCI University of California, Irvine | North Carolina Agricultural and Technical State University

Publications

1. Al-Shareeda, S., Ozguner, F., & Canberk, B. “Group-Signature Authentication to Secure Task Offloading in Vehicular Edge Twin Networks.” Workshop on BlockSecSDN, IEEE GLOBECOM 2024, Cape Town, South Africa, December 2024.
2. Al-Shareeda, S., Celik, Y., Bilgili, B., Al-Dubai, A., & Canberk, B. “Accurate AI-Driven Emergency Vehicle Location Tracking in Healthcare ITS’s Digital Twin.” IEEE MENACOMM 2025, Beirut, Lebanon, February 2025.
3. Al-Shareeda, S., Ozguner, F., Redmill, K., Duong, T. Q., & Canberk, B. “Lightweight Authenticated Task Offloading in 6G-Cloud Vehicular Twin Networks.” IEEE WCNC 2025, Milan, Italy, March 2025.
4. Al-Shareeda, S., Srinivasan, V., Al-Mudhaf, M., Bin Salamah, Y., Jabr, B., & Ozguner, F. “Evaluating Handover Impact on IBC-Authenticated Task Offloading in C-V2X VTENs.” IEEE VNC 2025, Portugal, June 2025.
5. Al-Shareeda, S., Saim, M., Jabr, B., Bin Salamah, Y., Alanazi, F., Yurdakul, G., Ozguner, F., & Ozguner, U. “When Pedestrians Hesitate: PPO-Based RL Collision Avoidance in Uncertain Scenarios.”. SmartNets 2025, Istanbul, July 2025.
6. Al-Shareeda, S., Saeed, N., Redmill, K., Jabr, B. A., Bin Salamah, Y., Al-Dubai, A., & Ozguner, F. “Novel DT-assisted Vehicular Task Offloading for Cloud, Edge, and Hybrid Deployments.” IEEE Transactions on Vehicular Technology (TVT), October 2025.

Miscellaneous

1. Awarded Best Runner-Up Paper: SmartNets 2025 Conference Presentation “When Pedestrians Hesitate: PPO-Based RL Collision Avoidance in Uncertain Scenarios.” Istanbul, July 22-24, 2025.
2. OSU CAR News Feature: “Evasive Maneuvers: New AI Training Program Ensures Safer Travel with Hesitant Pedestrians” September 2025.
3. Invited Seminar: Al-Shareeda, S. *“Shifting Ground: How AI Transforms Skills, Research, and the Competence Ecosystem.”* DevFest Basra 2025, organized by GDG Basra (Google Developers Group) and Women Techmakers Basra. Basrah, Iraq, December 2025.

Abstract

Vehicle-to-Everything (V2X) communication underpins Highly Automated Transportation Systems (HATS), yet scalability is constrained by three coupled challenges: (i) authentication overhead in latency-sensitive communication, (ii) mobility-aware computation offloading under dynamic wireless and infrastructure conditions, and (iii) safety-critical decision-making under behavioral uncertainty. Existing architectures often treat security, offloading, synchronization, and control as loosely coupled components, leading to systematic latency underestimation and degraded responsiveness in dense Cellular-V2X (C-V2X) environments.

This work advances an integrated cyber-layer intelligence framework evaluated through three coordinated tracks that explicitly couple authentication, mobility, synchronization, and learning within a unified decision loop.

Track 1 embeds Group Signatures (GS) and Identity-Based Cryptography (IBC) directly into end-to-end latency modeling and optimization. Results show that authentication introduces a persistent baseline delay and that C-V2X handovers amplify end-to-end latency by up to 47% while nearly doubling authentication frequency in dense regimes. By incorporating these costs into mobility-aware optimization, the framework prevents delay underestimation and stabilizes offloading decisions under scale and handover.

Track 2 extends Digital Twins (DTs) from passive mirrors to predictive coordination surfaces. Prediction-enabled synchronization reduces cyber-physical misalignment by approximately 88-94% and translates into measurable deployment-level gains, with DT-assisted cloud, edge, and hybrid configurations consistently reducing latency relative to reactive baselines across density, task-size, and transmission-rate scaling.

Track 3 addresses behavioral uncertainty using an uncertainty-aware Reinforcement Learning (RL) controller for autonomous safety. The proposed optimization policy achieves a collision rate of 0.0584 at a 10% hesitation probability and maintains bounded degradation as uncertainty increases, demonstrating adaptive safety without reliance on explicit trajectory prediction.

Overall, the results show that secure vehicular intelligence is inherently coupled: authentication and mobility define the latency structure, synchronization fidelity determines decision timeliness, and adaptive control governs safety under uncertainty. By coordinating these dimensions within a predictive cyber layer, the framework improves scalability, robustness, and safety in dense and dynamic vehicular environments, directly supporting USDOT and CARMEN+ priorities for secure and resilient intelligent transportation systems.

Note:
The results presented in this report are published in preferred academic conferences and journals.

Table of Contents

List of Figures

List of Tables

Table **Page**

Chapter 1
Introduction, Background, and Objectives

Vehicle-to-Everything (V2X) communication supports Highly Automated Transportation Systems (HATS) by enabling vehicles, infrastructure, and cloud or edge services to exchange safety-critical information in real time. As system density and automation levels increase, system scalability is increasingly constrained by three tightly coupled challenges: (i) authentication overhead in latency-sensitive communication, (ii) computation offloading under dynamic wireless and infrastructure conditions, and (iii) safety-critical decision-making under behavioral uncertainty. These challenges are not independent. Authentication mechanisms introduce signing, verification, and signaling costs that directly affect end-to-end delay. Under mobility, especially during Cellular-V2X (C-V2X) handovers, these costs can increase due to reauthentication and forwarding procedures. Offloading strategies that do not explicitly model authentication and mobility effects may therefore degrade in dense environments. In parallel, cyber-physical misalignment between physical vehicles and their digital representations can amplify delay and reduce responsiveness in safety-critical scenarios. Reactive control policies that ignore uncertainty further compound this effect, particularly in pedestrian interactions characterized by stochastic behavior.

To address these interdependent limitations, this report advances a unified cyber-layer intelligence framework in which communication, security, synchronization, and control are modeled as coordinated sequential decision processes. Rather than treating authentication, offloading, Digital Twin (DT) management, and safety control as isolated subsystems, the framework integrates them within a shared decision architecture driven by predictive modeling and Reinforcement Learning (RL)-based optimization. The framework is structured around three coordinated technical thrusts:

- **Track 1: Security-aware cyber-layer offloading under authentication and mobility.** Authenticated task offloading mechanisms are designed and evaluated in vehicular twin networks. The analysis quantifies authentication overhead, congestion effects, and handover-induced latency amplification, and incorporates these factors into adaptive decision-making models.

- **Track 2: DT architectures for predictive synchronization and optimization.** DT frameworks are enhanced with predictive models to reduce cyber-physical misalignment and with optimization-driven orchestration for cloud, edge, and hybrid deployments.

- **Track 3: RL-based safety under behavioral uncertainty.** An uncertainty-aware RL framework is developed for autonomous vehicle collision avoidance at unmarked crosswalks with stochastic pedestrian behavior.

Across all thrusts, the common design principle is cyber-layer decision intelligence, where communication, security, synchronization, and control are jointly modeled as state-aware processes and optimized using learning and predictive mechanisms, as depicted in Fig. 1.1. This architectural perspective positions the cyber layer not as a passive monitoring entity, but as an anticipatory optimization engine coordinating physical and computational resources under mobility and uncertainty.

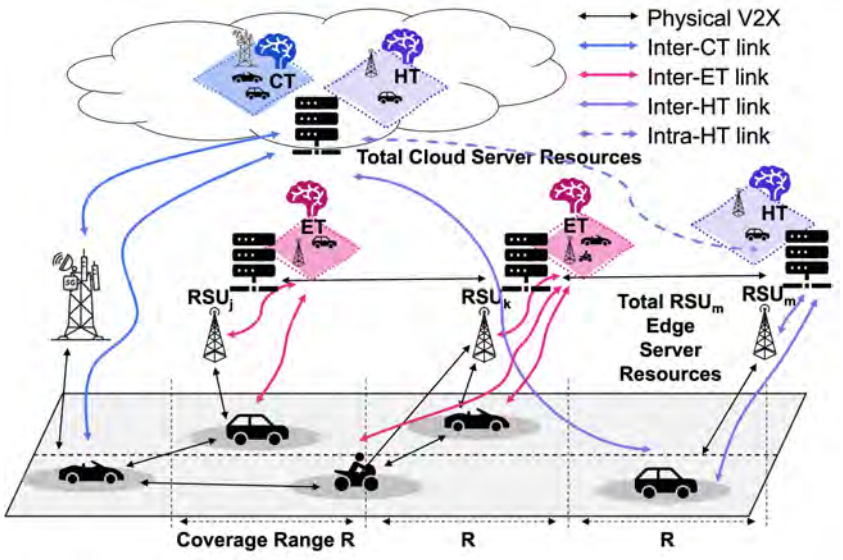

Figure 1.1. DT-assisted Cloud-Twin (CT), Edge-Twin (ET), and Hybrid-Twin (HT) deployment architectures. The cyber layer maintains synchronized DTs and executes RL-based optimization for mobility-aware offloading.

1.1 Background and Motivation

1.1.1 Authentication overhead in vehicular communication

Secure vehicular communication requires authentication mechanisms that provide integrity, accountability, and privacy preservation. These mechanisms introduce computation and communication overhead. In dense networks, frequent message exchange and topology changes increase signing and verification load, extending end-to-end latency [1–3]. Certificate-based approaches further add transmission and validation overhead due to certificate exchange and revocation handling. Mobility increases this effect. During handovers, vehicles may undergo repeated authentication and signaling exchanges, increasing delay and infrastructure load [4–6]. As network density and task size increase, authentication overhead becomes a measurable performance component rather than a negligible constant. This project evaluates two authentication approaches suitable for vehicular twin settings. Group Signatures (GS) provide anonymity with conditional traceability but incur signing and verification costs that must be incorporated into latency models [7–9]. Identity-Based Cryptography (IBC) removes certificate exchange but still requires explicit modeling of signing, verification, and key-management assumptions [1, 10]. Both mechanisms are evaluated within mobility-aware offloading workflows in vehicular twin architectures [11, 12].

1.1.2 Computation offloading and cyber-layer coordination

Connected vehicles rely on computation offloading to meet stringent latency and processing requirements [13–15]. Edge-based processing reduces round-trip delay relative to centralized cloud execution, while cloud infrastructure offers higher aggregate computational capacity [16, 17]. Hybrid cloud-edge models attempt to balance these characteristics through dynamic resource allocation across heterogeneous layers [4, 18]. Offloading performance depends on wireless channel conditions, infrastructure load, task heterogeneity, and mobility [19, 20]. Traditional approaches often rely on reactive heuristics or simplified delay estimators, limiting scalability and coordination under dynamic conditions [21, 22].

DT architectures provide synchronized cyber representations of vehicles, infrastructure, and communication links [11, 12, 23]. These representations enable state-aware latency estimation and proactive decision-making by integrating mobility, communication state, and compute availability into unified optimization frameworks [18, 24]. DT-assisted offloading therefore allows mobility, link conditions, and resource constraints to be incorporated directly into learning-based optimization loops.

1.1.3 Safety under behavioral uncertainty

Pedestrian interaction at unmarked crosswalks presents uncertainty due to hesitation, reversals, and context-dependent decision patterns. Behavioral variability is influenced by environmental conditions, social context, and perceived risk. Hesitation and crossing confidence have been modeled using observable signals such as gaze behavior, speed, and time-to-collision [25]. Game-theoretic formulations represent vehicle-pedestrian interaction as negotiation under partial observability [26, 27]. Additional studies highlight non-rational and environment-dependent pedestrian responses [28]. Broader empirical analyses further examine contextual and interaction-driven effects on crossing decisions [29–34]. RL provides a data-driven framework for adaptive decision-making without relying on fixed behavioral assumptions. Prior work considers pedestrian-as-agent formulations [35–38], AV-as-agent formulations [39–43], and multi-agent RL approaches [44, 45]. However, practical deployment constraints, such as limited observability, onboard computation limits, and real-time decision requirements, motivate AV-centered learning with stochastic pedestrian modeling. This project adopts that direction and evaluates safety performance under varying hesitation probabilities using uncertainty-aware policy learning.

Collectively, authentication overhead modeling, mobility-aware offloading, predictive DT synchronization, and uncertainty-aware control represent tightly coupled system dimensions. Authentication costs influence offloading latency; offloading delay affects synchronization fidelity; synchronization lag impacts safety-critical responsiveness. Modeling these effects independently obscures their interaction.

By embedding them within a coordinated cyber-layer decision framework, this project establishes a unified architecture that explicitly captures their interdependence and enables anticipatory optimization under dense and dynamic vehicular conditions.

1.2 Project Objectives

This project develops a security-aware and learning-driven cyber-physical framework for next-generation vehicular networks. The framework integrates authenticated task offloading, DT synchronization, and RL-based safety control within a unified system. The objectives are designed to address authentication overhead, mobility-aware optimization, predictive synchronization, and uncertainty-aware autonomy, while aligning with USDOT and CARMEN+ priorities on resilient infrastructure and secure intelligent transportation systems.

1.2.1 Objective 1: Design Security-Aware Offloading Architectures

The first objective is to develop authenticated task offloading mechanisms in vehicular twin networks where security operations are directly integrated into latency-sensitive performance models. Instead of treating authentication as a separate or external step, this objective embeds cryptographic signing and verification costs into end-to-end delay formulations. To achieve this, the project designs GS-based authenticated offloading mechanisms for Vehicular Edge Twin Networks (VETNs) and IBC frameworks for cloud-enabled vehicular twin deployments. The models explicitly capture authentication behavior under C-V2X mobility and handover events, and quantify signing and verification overhead as part of total task latency. By incorporating cryptographic cost into system modeling, this objective establishes secure V2X communication architectures that resist impersonation and replay attacks while preserving anonymity and traceability under dense and mobile conditions.

1.2.2 Objective 2: Develop RL-Based Intelligent Offloading Optimization

The second objective is to formulate authenticated task offloading as a sequential decision problem and develop a learning-based optimization engine that adapts to changing network conditions. Rather than relying on static delay estimation, this objective enables adaptive cyber-layer orchestration. Cloud, edge, and hybrid deployment modes are modeled as Markov Decision Processes (MDPs). The learning agent jointly optimizes binary offloading decisions and continuous resource allocation while accounting for mobility dynamics, congestion, and authentication overhead. The overall goal is to minimize authenticated task latency under computational and deadline constraints. This objective enables adaptive edge-cloud coordination and learning-driven infrastructure optimization in dynamic vehicular environments.

1.2.3 Objective 3: Establish DT-Enabled Predictive Synchronization

The third objective focuses on improving cyber-physical consistency by extending DT architectures with predictive synchronization capabilities. Instead of passively mirroring system state, the DT layer is enhanced to forecast state evolution and reduce synchronization delay. The project models synchronization gaps between physical vehicles and their cyber replicas and implements spatiotemporal prediction models using Support Vector Regression (SVR) and Deep Neural Networks (DNNs). These models reduce Digital Twin lag ($\Delta\tau$) through forward state estimation, enabling proactive decision-making in the cyber layer. This objective strengthens timing accuracy and coordination in safety-critical ITS deployments, including emergency healthcare mobility networks.

1.2.4 Objective 4: Implement and Evaluate DT-Assisted Cloud-Edge-Hybrid Deployments

The fourth objective is to design and evaluate DT-assisted deployment configurations that integrate authentication modeling, mobility-aware latency analysis, and learning-based optimization within unified cloud, edge, and hybrid architectures. The project implements Cloud-Twin (CT), Edge-Twin (ET),

and Hybrid-Twin (HT) deployment modes and compares them against traditional reactive baselines. Handover-aware latency modeling is incorporated to capture mobility-driven delay amplification, and wireless channel variations are integrated into cyber-layer decision-making. This objective provides a scalable evaluation framework for resilient and mobility-aware infrastructure deployment.

1.2.5 Objective 5: Enable RL-Driven Safety under Behavioral Uncertainty

The fifth objective extends RL beyond infrastructure optimization into safety-critical vehicle control under uncertain pedestrian behavior. This objective addresses decision-making in environments characterized by stochastic hesitation and partial observability. Pedestrian hesitation in unmarked crosswalk scenarios is modeled as a probabilistic process, and autonomous vehicle control is formulated as an MDP. RL agents are trained to balance collision avoidance, efficiency, and socially compliant behavior. Collision risk is quantified using Time-To-Collision (TTC)-aware reward structures. This objective supports uncertainty-aware autonomy and safety validation in mixed-traffic environments, consistent with USDOT priorities on autonomous vehicle safety assurance.

Chapter 2

Methodology: Integrated Cyber-Physical System Architecture and Optimization

This chapter formalizes the unified cyber-layer intelligence architecture introduced in Chapter 1. The framework addresses three tightly coupled challenges in next-generation vehicular networks: authentication overhead in latency-sensitive communication [1, 2], mobility-aware computation offloading under dynamic infrastructure conditions [19, 20], and safety-critical decision-making under behavioral uncertainty [39, 44]. Conventional vehicular systems treat these challenges independently. Authentication is implemented as an external security layer; offloading is optimized without explicit modeling of mobility-aware security overhead; and safety control is designed independently from infrastructure coordination. Such compartmentalization obscures the interaction between security cost, mobility dynamics, synchronization delay, and control responsiveness. In contrast, the proposed framework models communication, computation, and control as interdependent components within a unified cyber-physical decision architecture. Security-aware task offloading, DT synchronization [11, 23], mobility-aware latency modeling [4], and RL-based optimization are integrated into a closed-loop coordination framework. The cyber layer functions as an anticipatory decision engine that continuously evaluates authenticated execution cost, infrastructure variability, and vehicular mobility to enable predictive and adaptive optimization.

2.1 Architecture Overview and Closed-Loop Operation

The proposed system, Fig. 1.1 and Fig. 2.1, operates as a layered cyber-physical architecture in which physical entities and cyber-layer intelligence exchange state information and control decisions

in continuous feedback. The design follows emerging vehicular DT paradigms [11, 18], but extends them by embedding authentication cost modeling and learning-based optimization directly within the cyber layer. At the physical layer, vehicles generate computational tasks and update mobility states. These states are mirrored in the cyber layer through synchronized DT replicas. The cyber layer evaluates expected execution latency under CT, ET, and HT deployment configurations while explicitly accounting for authentication overhead, transmission variability, and handover risk. An RL-based optimization engine selects offloading and resource allocation decisions. These decisions are executed in the physical layer, and the resulting performance metrics update DT states, closing the cyber-physical feedback loop. This iterative structure transforms offloading from a reactive scheduling procedure into

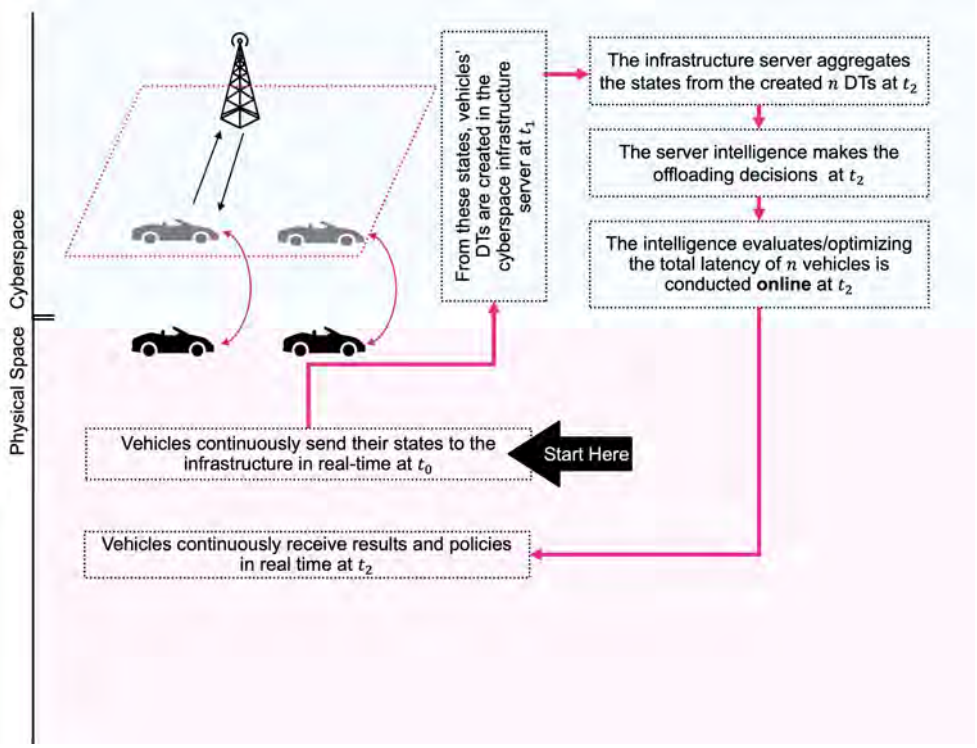

Figure 2.1. Closed-loop cyber-physical operation: synchronized DT state acquisition, authentication-aware latency evaluation, RL-based decision selection, physical execution, and state update.

a predictive decision process conditioned on authentication cost, infrastructure capacity, and mobility dynamics.

2.2 System Design

2.2.1 Physical Layer Modeling

The physical layer consists of n mobile vehicles, m Roadside Units (RSUs), and a centralized cloud server, consistent with hierarchical vehicular edge-cloud architectures [16, 17]. Each vehicle possesses onboard computational capacity f_{V_i}, each RSU provides aggregate computational capacity F_{RSU_j}, and the cloud server provides computational capacity F_C. Wireless communication occurs over Vehicle-to-Infrastructure (V2I) and cellular Vehicle-to-Cloud (V2C) links whose transmission rates vary dynamically according to channel conditions and infrastructure load [4, 5]. Vehicles generate authenticated safety-critical tasks defined as:

$$S_{V_i} = \{d_{V_i}, c_{V_i}, \tau_{V_i}, \bar{d}_{V_i}\} \tag{2.1}$$

where d_{V_i} denotes input data size, c_{V_i} required CPU cycles per byte, τ_{V_i} task deadline, and $\bar{d}_{V_i}$ result size. The end-to-end latency is influenced by communication delay, queueing delay, computation time, authentication overhead, and potential handover effects. These components are jointly modeled to capture realistic execution behavior under dynamic vehicular mobility.

2.2.2 Security Layer Integration

Authentication is embedded directly into the offloading workflow through two mechanisms: GS-based anonymous authentication for VETNs [7, 9], and IBC-based authentication for 6G-enabled cloud vehicular twin deployments [1, 10]. Unlike conventional approaches that treat authentication as an external verification stage [3], the proposed framework incorporates cryptographic overhead explicitly into latency modeling. Let T_{sign} and T_{verify} denote signing and verification costs. The resulting end-to-end latency model is defined in Chapter 3, Section 3.3, where all latency components are instantiated consistently across deployments. $T_{handover}$ captures additional delay caused by mobility-induced coverage transitions and reauthentication procedures. The security layer guarantees existential unforgeability under chosen message attacks, replay resistance via timestamp validation, anonymity preservation,

and conditional traceability through trusted authorities. Embedding authentication cost within performance models enables quantitative tradeoff evaluation between security strength and latency efficiency in dense vehicular networks. Fig. 2.2 and Fig. 2.3 exhibit a GS- and IBC-based security layer integration in offloading.

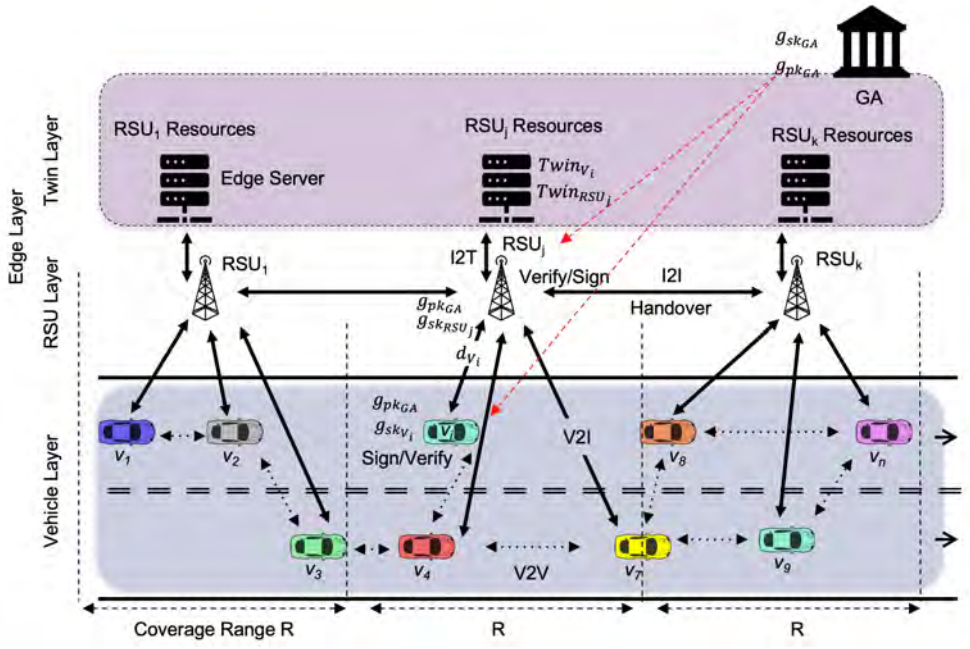

Figure 2.2. GS-based Security layer integration in offloading.

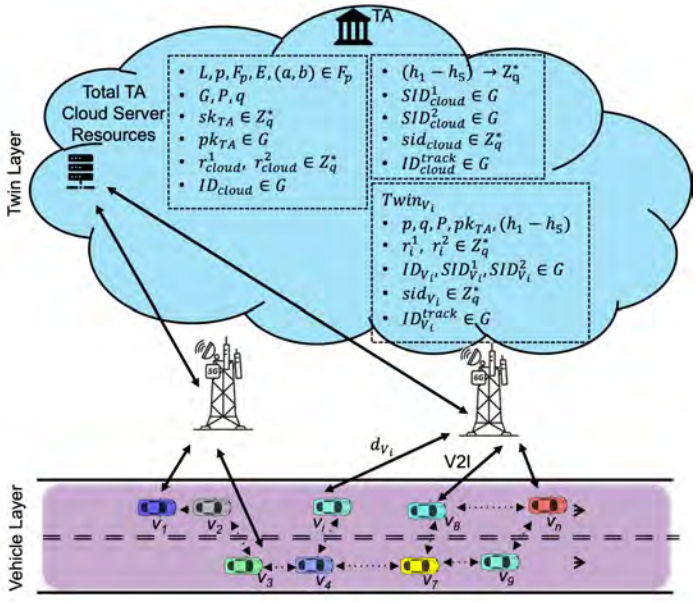

Figure 2.3. IBC-based Security layer integration in offloading.

2.2.3 Digital Twin Layer

Each physical entity maintains a synchronized DT replica [11, 23]. The DT of vehicle V_i at time t, denoted $DT_{V_i}^t$, maintains mobility state, task parameters, and computational attributes. Similarly, $DT_{RSU_j}^t$ represents infrastructure capacity and coverage, DT_C^t cloud availability, and DT_{link}^t instantaneous uplink and downlink transmission rates. DT synchronization occurs at discrete intervals Δt, enabling forward estimation of system state and predictive evaluation of execution outcomes [24]. The synchronized DT state forms the observable space for the learning-based optimization framework. By maintaining mobility-aware and authentication-aware state representations, the DT layer enables proactive orchestration of vehicular computation rather than reactive scheduling.

2.3 Mobility and Handover Modeling

Vehicle mobility directly influences both latency modeling and state transitions. A handover event occurs when a vehicle exits RSU coverage during task execution:

$$y_{V_i} = \begin{cases} 0 & v_{V_i} \cdot T_{offload} \leq R \\ 1 & \text{otherwise} \end{cases} \tag{2.2}$$

where v_{V_i} denotes vehicle speed, $T_{offload}$ expected offloading duration, and R RSU coverage radius. When $y_{V_i} = 1$, additional latency is incurred due to backhaul transmission, reauthentication cost, and result forwarding delay. This mobility indicator directly influences both T_{total} and the transition dynamics of the optimization model, tightly coupling physical-layer movement with cyber-layer decision-making as shown in Fig. 2.4.

2.4 RL-Based Optimization Framework

Here, the authenticated task offloading is formulated as an MDP $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$. The state space $\mathcal{S}$ is constructed from synchronized DT variables, including task size, link rates, authentication overhead, mobility state, and computational capacity as in Fig. 2.5. The action space $\mathcal{A}$ includes binary offloading decision $X_{V_i} \in \{0, 1\}$ and continuous allocation variables $F_{RSU_j}^{V_i}$ and $F_C^{V_i}$. The reward function is defined

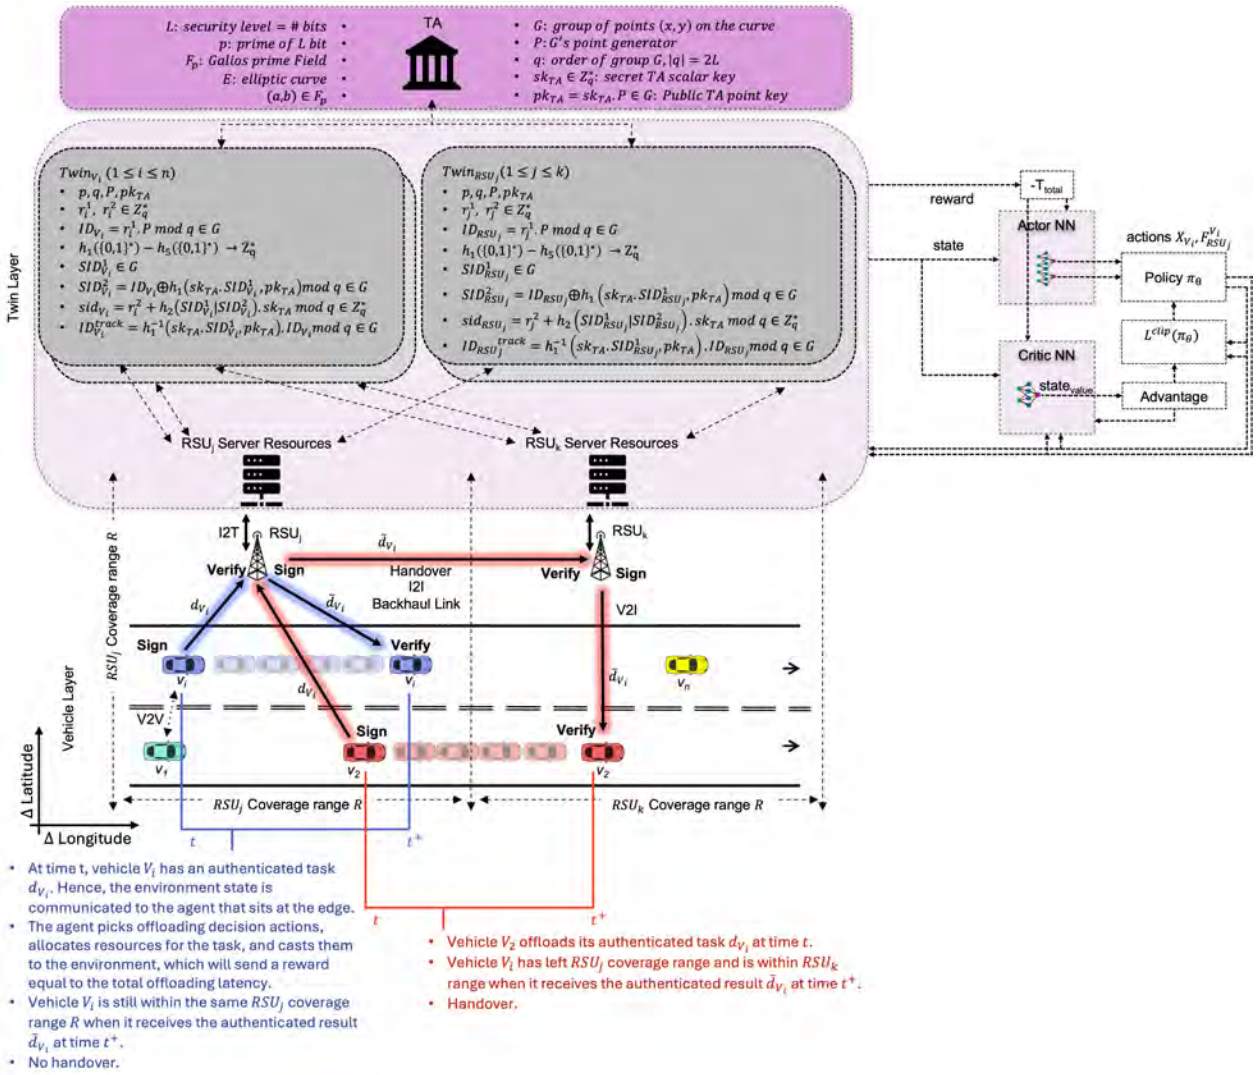

Figure 2.4. IBC-based Security layer integration in offloading with/without handover.

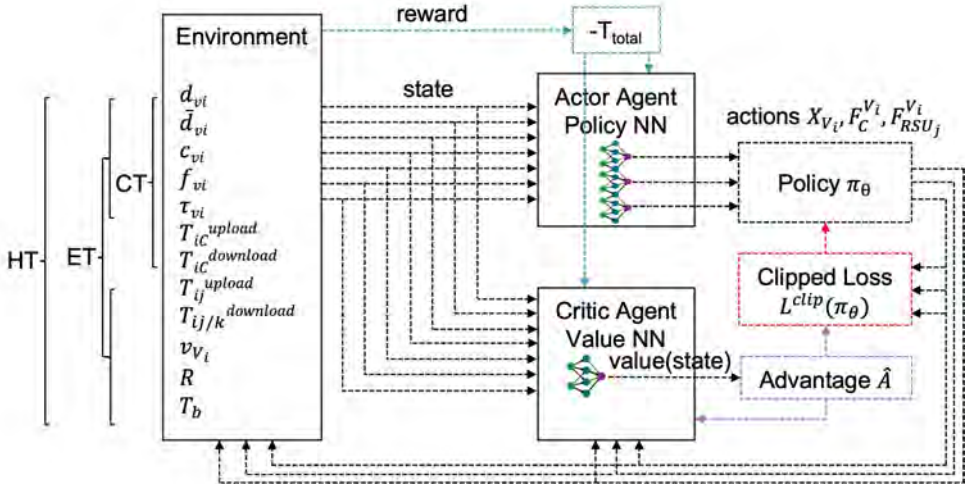

Figure 2.5. Used PPO-based RL Framework.

as:

$$R_t = -T_{total}(t) \tag{2.3}$$

enforcing minimization of authenticated latency. The optimization engine adopts a Proximal Policy Optimization (PPO) actor-critic architecture with two hidden layers using ReLU activation. Dual output branches generate categorical offloading decisions and Gaussian allocation parameters. Policy updates follow:

$$L^{PPO} = L^{clip} - H_{coeff} \cdot H(\pi_\theta) \tag{2.4}$$

where L^{clip} ensures stable updates and H_{coeff} regulates exploration. The agent operates entirely within the cyber layer using synchronized DT states, enabling predictive mobility-aware optimization.

2.5 Safety-Critical RL Framework

Extending the same cyber-layer decision paradigm to vehicle control, an RL framework addresses autonomous safety under stochastic pedestrian behavior. The environment state includes vehicle and pedestrian positions, velocities, crosswalk geometry, and TTC. The autonomous vehicle selects among four discrete maneuvers: maintain speed, brake, lateral avoidance left, or lateral avoidance right as in Fig. 2.6. The reward function balances collision avoidance, TTC maximization, efficiency preservation, and socially compliant behavior. Although this safety control agent operates on a different control surface than the infrastructure-level offloading agent, both modules share the same architectural principle: synchronized state acquisition, predictive consequence evaluation, and performance-driven policy optimization under uncertainty. This unified cyber-layer design ensures coherence between infrastructure coordination and safety-critical autonomy within the proposed vehicular computing framework.

The proposed architecture integrates authentication modeling, mobility-aware latency analysis, DT synchronization, and RL–based optimization within a unified closed-loop cyber-physical framework. With the system formulation, state representations, and optimization mechanisms established, the following chapter presents the experimental setup used to quantitatively validate the design.

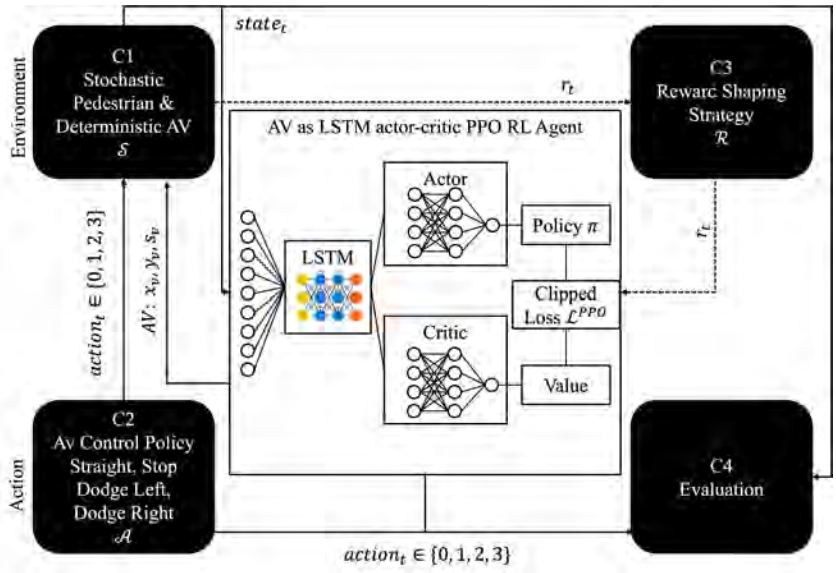

Figure 2.6. Safety-Critical RL Framework.

Chapter 3

Experimental Setup and Evaluation Framework

This chapter presents the unified experimental framework used to evaluate the secure offloading architectures, DT synchronization mechanisms, and RL-based safety models developed in this work. The evaluation is designed to quantify authentication overhead, mobility-induced delay amplification, DT synchronization fidelity and lag reduction, and RL learning stability under controlled and reproducible conditions. The experimental environment mirrors the integrated cyber-physical architecture described in Chapter 2. Secure offloading, mobility modeling, DT synchronization, and RL-based optimization are implemented within a single simulation pipeline to enable coherent system-level analysis across consistent assumptions and shared state representations.

3.1 Simulation Environment

All experiments are conducted using a custom-built simulation framework implemented in Python and PyTorch. The framework integrates (i) vehicular mobility and handover modeling, (ii) wireless communication and backhaul delay modeling, (iii) task generation and queueing dynamics, (iv) explicit authentication overhead injection, (v) DT synchronization, and (vi) RL decision engines within a unified executable environment.

3.1.1 Mobility and Handover Trace Generation

Vehicular mobility is modeled using stochastic trajectory generation to represent density variation and C-V2X handover events. Vehicles move through RSU coverage regions; a handover event is triggered when a vehicle exits the serving RSU coverage while a task is still in transmission or execution. Mobility

traces are generated with controlled random seeds to ensure reproducibility, and are parameterized to support multiple density regimes and speed profiles as required by the studied scenarios.

3.1.2 Wireless and Backhaul Link Modeling

Wireless access links are modeled with Rayleigh fading to capture stochastic channel variation under mobility. The instantaneous transmission behavior is reflected through time-varying effective rates used in the uplink and downlink latency terms. Backhaul links are explicitly modeled to account for RSU-to-RSU and RSU-to-cloud forwarding delays that occur during handovers and in hybrid execution paths.

3.1.3 Task and Queueing Model

Task generation incorporates heterogeneous input sizes, computational requirements (e.g., CPU cycles), and deadline constraints to emulate latency-sensitive vehicular workloads. For edge and cloud execution, queueing dynamics are incorporated through per-server queues and service rates. The resulting queueing delay term is included in the end-to-end latency model in Section 3.3.

3.1.4 Authentication Overhead Injection

Authentication overhead is explicitly modeled for both GS and IBC mechanisms. Cryptographic signing and verification costs are injected directly into the latency computation pipeline. Under mobility-induced handovers, reauthentication events are triggered and accounted for as additional overhead. All cryptographic parameters used to compute signing and verification costs are logged for validation and are consistent with the authentication models defined earlier.

3.1.5 DT Synchronization and Misalignment

For DT-assisted deployments, synchronized cyber replicas of vehicles, RSUs, cloud servers, and communication links are maintained at predefined synchronization intervals. The cyber layer uses synchronized DT states to support mobility-aware decisions and predictive coordination. In contrast,

traditional baselines rely on reactive state estimation without predictive DT coordination. Fig. 3.1

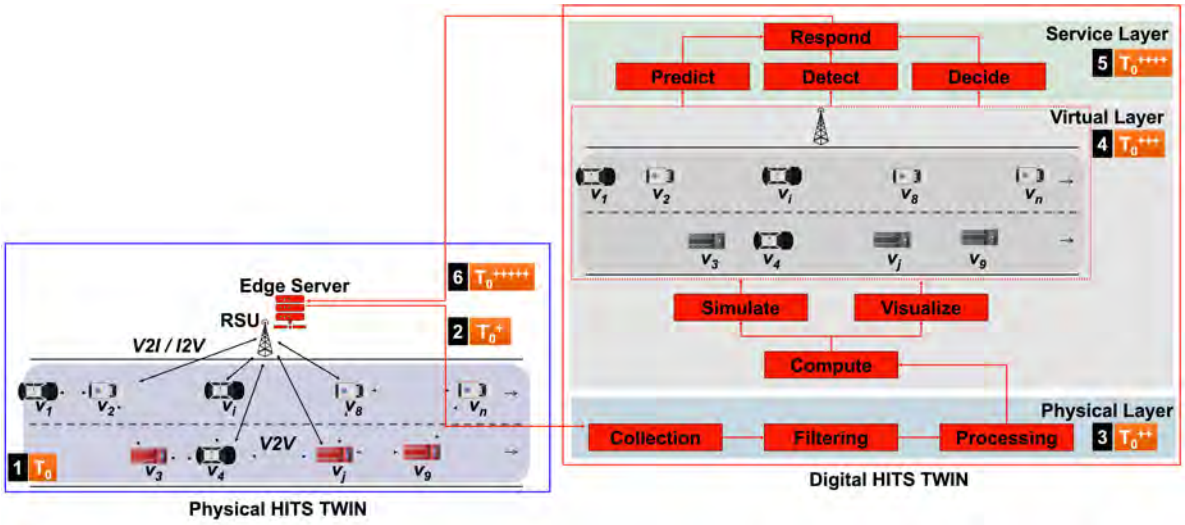

Figure 3.1. Illustration of temporal misalignment between physical entities and DT representations. Predictive synchronization reduces the lag $\Delta\tau$ between physical state updates and cyber-layer decisions.

illustrates the synchronization gap between physical entities and their cyber replicas. This temporal misalignment motivates the DT lag metric $\Delta\tau$ and its reduction under predictive synchronization, as defined in Section 3.5.

3.2 Deployment Configurations

To isolate the impact of DT synchronization and learning-based optimization, six deployment configurations are evaluated. DT-assisted configurations include: CT, ET, and HT. In these deployments, the cyber layer maintains synchronized DT state awareness and executes learning-based optimization for task offloading and resource allocation. For comparison, three traditional baselines are implemented:

- Traditional Cloud (TC),
- Traditional Edge (TE),
- Traditional Hybrid (TH).

Traditional deployments operate reactively using static delay estimation without predictive DT coordination or adaptive learning. Fig. 1.1 presents the DT-assisted architecture used in CT, ET, and HT

modes. The HT configuration dynamically selects between cloud and edge execution paths based on predicted mobility and latency conditions.

3.3 Latency and Mobility Modeling

End-to-end task latency is modeled as the sum of transmission, queueing, computation, cryptographic overhead, and mobility-induced handover delay:

$$T_{\text{total}} = T_{\text{tx}} + T_{\text{queue}} + T_{\text{comp}} + T_{\text{sign}} + T_{\text{verify}} + T_{\text{handover}}. \tag{3.1}$$

Here, T_{sign} and T_{verify} capture cryptographic signing and verification overhead, while T_{handover} represents mobility-induced delay. T_{tx} includes uplink and downlink components across the selected execution path (edge, cloud, or hybrid), using the access and backhaul link models described in Section 3.1.2. T_{queue} captures waiting time at edge RSUs and cloud servers under task arrivals and service rates. T_{comp} captures execution time under allocated compute capacity and task computational demand. A handover event is triggered when a vehicle exits RSU coverage during task execution or result delivery. The handover delay term is decomposed as:

$$T_{\text{handover}} = T_{\text{bh,fw}} + T_{\text{reauth}} + T_{\text{result,fw}}, \tag{3.2}$$

where $T_{\text{bh,fw}}$ denotes backhaul forwarding delay (e.g., RSU-to-RSU forwarding), T_{reauth} denotes reauthentication delay (including any required cryptographic operations and signaling), and $T_{\text{result,fw}}$ denotes forwarding delay for task results to reach the vehicle under mobility. This formulation enables measurement of mobility-driven latency amplification under varying speed and density conditions, and supports direct comparison between DT-assisted and traditional deployments under identical mobility traces.

3.4 RL Training Configuration

Two RL agents are evaluated under this unified framework:

- A PPO-based offloading optimization agent for CT, ET, and HT deployments.

- A PPO-LSTM-based safety control agent for pedestrian interaction scenarios.

3.4.1 PPO Agent for DT-Assisted Offloading

The PPO agent is implemented using an actor-critic architecture with two hidden layers and ReLU activation. Dual output branches produce (i) categorical offloading decisions and (ii) continuous resource allocation parameters (modeled via a Gaussian distribution). The agent operates on synchronized DT states in CT, ET, and HT, and its reward is defined using latency and constraint satisfaction terms consistent with the problem formulation in the secure offloading chapters. Training uses iterative PPO updates with clipped policy optimization. The configuration parameters (iterations, episodes, environment steps, learning rate, discount factor, entropy coefficient, and PPO clip) are logged and fixed per experiment to ensure repeatability. The following training configuration is used for the DT-assisted offloading experiments in this report:

- Training iterations: 10,000,
- Episodes per iteration: 100,
- Environment steps per episode: 100,
- Learning rate: 0.003,
- Discount factor: $\gamma = 0.9$,
- Entropy coefficient: 0.08.

Convergence is evaluated using cumulative reward stabilization and variance analysis across independent training runs (Section 3.5). Fig. 2.5 illustrates the structured PPO-DT training process. The agent samples actions from synchronized DT states, executes offloading decisions, computes rewards based on T_{total} in (3.1) and related constraints, and updates policy parameters using the clipped PPO objective.

3.4.2 PPO-LSTM Agent for Safety Control

The safety model addresses autonomous vehicle navigation under stochastic pedestrian hesitation at unmarked crosswalks. The state space includes vehicle and pedestrian positions, velocities, crosswalk geometry, and TTC. The action space consists of four discrete maneuvers: maintain speed, brake, dodge left, and dodge right. The policy is optimized using PPO augmented with an LSTM module to capture temporal dependencies induced by pedestrian hesitation dynamics. The reward function balances collision avoidance, TTC maximization, efficiency preservation, and socially compliant behavior. The safety agent training and evaluation follow the configuration and scenario definitions provided in the safety chapter and shown in Fig. 2.6, while sharing the reproducibility and validation protocol described in Section 3.6.

3.5 Performance Metrics

Performance metrics are selected to quantify system-level impact across security, synchronization, mobility, and safety dimensions. Let $\mathcal{E}$ denote the set of evaluation episodes (or tasks) under a given configuration. For **secure offloading experiments**, the following metrics are measured:

- Total system latency (ms): T_{total} in (3.1).
- Deadline satisfaction ratio: the fraction of tasks satisfying $T_{\text{total}} \leq D$, where D is the task deadline.
- Resource utilization efficiency: utilization of edge and cloud compute resources under the selected allocation policy.
- Authentication frequency under mobility: frequency of authentication and reauthentication events per task or per handover.
- Handover-induced delay amplification: relative increase in T_{total} attributable to T_{handover} in (3.2).

For **DT evaluation**, synchronization fidelity and DT lag reduction are measured. Synchronization fidelity quantifies the consistency between physical states and DT states at the cyber layer under the selected synchronization interval, and is assessed using the DT lag metric $\Delta\tau$ and its statistical dispersion over time (e.g., variance or confidence interval). The DT lag $\Delta\tau$ captures the temporal misalignment between physical state updates and the DT state used by the cyber layer, as conceptually illustrated in Fig. 3.1. A reduction in $\Delta\tau$ quantifies predictive synchronization gains.

For **RL models**, learning stability and behavioral robustness are evaluated using:

- Cumulative reward progression: episode return trends over training updates.
- Collision rate (safety): collision events per episode under each ϕ_h setting.
- Episode length: number of steps until terminal condition (collision or successful completion).
- Action distribution: empirical selection frequency of each action to characterize learned behavior.

3.6 Reproducibility and Validation

To ensure reproducibility, all simulations are executed under documented configurations with fixed random seeds. Multiple seeds are used for RL experiments to evaluate variance sensitivity, and reported metrics are aggregated across independent runs using consistent aggregation rules (mean and confidence intervals when applicable). Hyperparameter sensitivity analysis is performed to confirm training stability under moderate parameter perturbations. All cryptographic cost parameters, mobility traces, network configurations, and training hyperparameters are logged and archived to enable independent validation.

3.7 Scope of Quantitative Evaluation

The unified evaluation framework enables systematic quantification of:

- Authentication overhead under scale and varying workload intensity,

- Mobility-induced latency amplification under handover events,
- Predictive synchronization impact through $\Delta\tau$ reduction,
- DT-assisted versus traditional deployment performance under matched conditions,
- RL-driven safety performance under controlled behavioral uncertainty.

By integrating authentication, mobility, synchronization, and learning within a single controlled environment, the evaluation reflects scalable, mobility-aware, and security-constrained vehicular intelligence under realistic operating conditions.

Chapter 4

Results and Analysis of Our Cyber-Layer Intelligence

This chapter consolidates the quantitative evaluation outcomes of the proposed cyber-layer intelligence framework under the unified experimental setup in Chapter 3. Results are reported using the same track structure established in Chapter 1. Across the three tracks, the common objective is to validate that cyber-layer decision intelligence yields measurable benefits in latency, robustness, and safety when authentication, mobility, DT synchronization, and control are treated as coupled system dimensions rather than isolated subsystems.

4.1 Security-Aware Offloading Under Authentication and Mobility

Track 1 evaluates secure task offloading under explicit cryptographic cost modeling and mobility-induced handover, following the total latency decomposition in (3.1). The key contribution is that authentication overhead is not a negligible constant in dense or mobile settings, and its interaction with handover procedures can measurably amplify end-to-end delay. By embedding signing and verification costs directly into the execution pipeline, the proposed framework enables latency-accurate decision-making and realistic scalability assessment for authenticated vehicular computing.

4.1.1 Authentication Overhead as a First-Order Latency Component

In the unified latency model, authentication enters as:

$$T_{\text{auth}} = T_{\text{sign}} + T_{\text{verify}}, \tag{4.1}$$

which directly increases T_{total} in (3.1). When tasks are latency-sensitive, repeated, or executed under high traffic density, T_{auth} becomes a dominant contributor and can shift the optimal offloading decision boundary between edge, cloud, and hybrid execution. This effect is consistent with the motivation presented in Chapter 1, Section 1.1.1, and is operationally captured in the GS and IBC integration pipelines shown in Fig. 2.2 and Fig. 2.3. We evaluate GS- and IBC-authenticated task offloading under varying densities, task sizes, transmission rates, and mobility conditions to establish the authentication-latency relationship and its scalability implications.

GS-authenticated offloading in VETNs (edge): Fig. 4.1 and Fig. 4.2 characterize the latency growth as vehicle density increases and quantify the recovery achievable via transmission scaling (edge access and inter-RSU backhaul). These results show that even when strong privacy is enforced using GS, latency can be partially mitigated through communication resource upgrades, with larger gains in sparse-to-moderate densities and diminishing returns as contention increases.

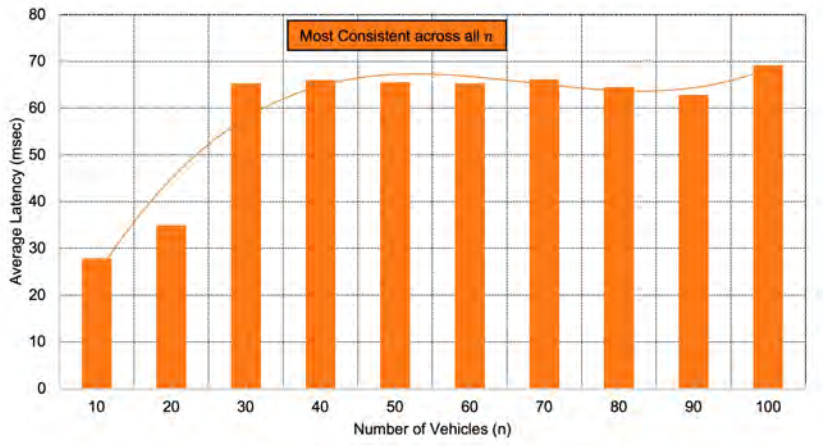

Figure 4.1. Average latency (msec) vs. number of vehicles n for GS-authenticated offloading in VETNs.

IBC-authenticated offloading in 6G-VTNs (cloud): Fig. 4.3 and Table 4.1 jointly quantify the authentication overhead impact under different task sizes, network scales, and transmission rates. The latency curves show that increasing the uplink and downlink rate consistently reduces delay, while the offloading percentage table reveals that authentication overhead can substantially depress offloading

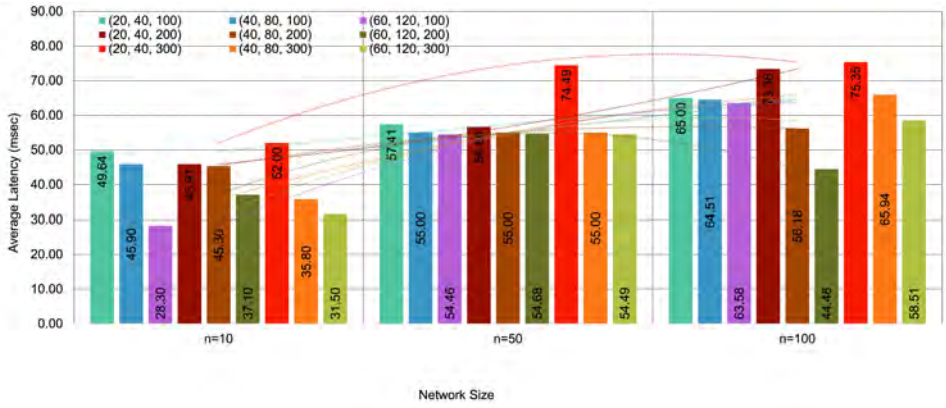

Figure 4.2. Performance vs. data rates at various network sizes for GS-authenticated offloading.

efficiency, particularly as density and task size grow. Transmission scaling partially restores feasibility but does not fully eliminate the scale sensitivity.

TABLE 4.1
OFFLOADING PERCENTAGE AT 100 MBPS, 500 MBPS, AND 1000 MBPS FOR IBC-AUTHENTICATED CLOUD OFFLOADING.

	100 Mbps Data Rate						500 Mbps Data Rate						1000 Mbps Data Rate					
	Task Size (B)						Task Size (B)						Task Size (B)					
	50 B		30 KB		300 KB		50 B		30 KB		300 KB		50 B		30 KB		300 KB	
n	(w/o IBC)	(w IBC)	(w/o)	(w)	(w/o)	(w)	(w/o)	(w)	(w/o)	(w)	(w/o)	(w)	(w/o)	(w)	(w/o)	(w)	(w/o)	(w)
10	9.73	4.13	8.13	3.65	6.35	3.25	6.30	2.00	15.80	12.03	11.85	8.55	17.18	15.65	15.65	9.10	12.85	9.55
40	3.58	1.73	3.72	2.97	4.22	1.57	5.18	3.56	3.56	1.65	3.35	2.88	3.16	2.67	2.67	1.74	2.60	1.63
70	2.00	0.55	2.31	1.80	1.00	0.80	2.65	1.35	2.41	1.49	1.42	1.24	2.69	1.12	1.58	1.12	2.31	0.00
100	1.00	0.58	0.80	0.68	0.53	0.50	1.36	0.00	1.24	0.51	1.21	0.56	1.01	0.24	1.44	1.24	1.44	0.56

4.1.2 Mobility and Handover Amplification Under Authentication

Mobility amplification is quantified through the handover decomposition in (3.2). This decomposition reveals that mobility introduces three additive delay components beyond baseline transmission and computation. Critically, authentication transforms handover from a pure forwarding event into a compound cryptographic and signaling event. As a result, T_{total} becomes structurally sensitive to coverage transitions. Table 4.2 quantifies average latency under no-handover and handover conditions across network densities and task sizes. Several observations are immediate:

- Latency increases monotonically with density under both conditions.
- The handover gap widens with task size, particularly for 300 KB and 3 MB workloads.
- The authentication penalty is magnified under mobility, especially at $n = 70$ and $n = 100$.

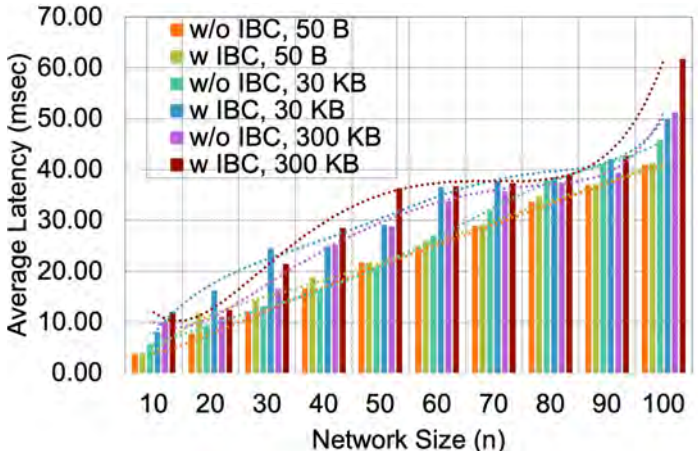

(a)

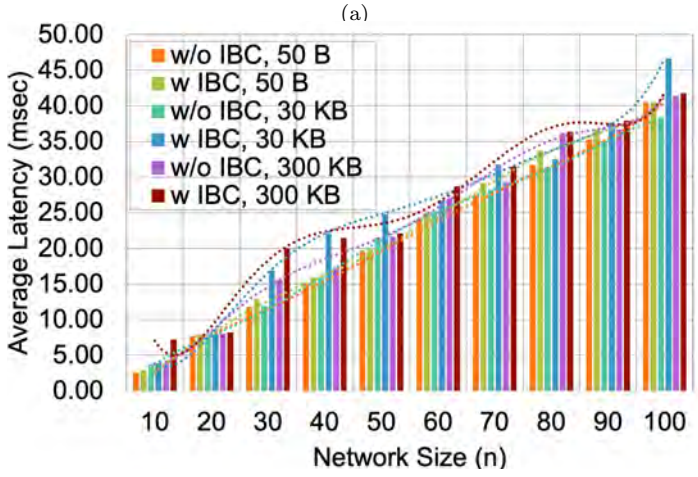

(b)

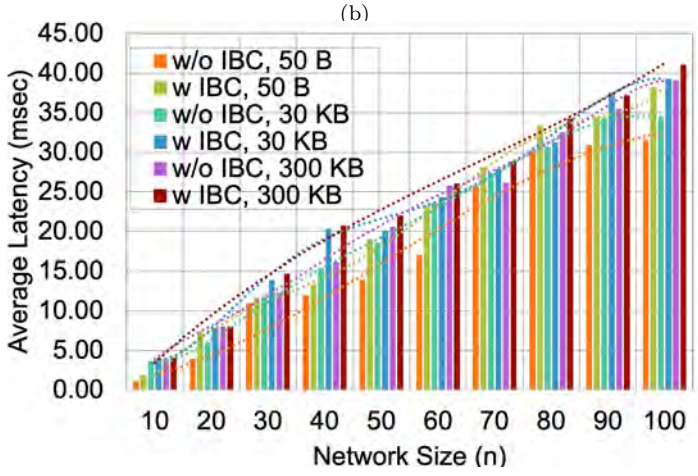

(c)

Figure 4.3. Average latency (msec) at different data rates for IBC-authenticated cloud offloading. (a) 100 Mbps. (b) 500 Mbps. (c) 1000 Mbps.

TABLE 4.2
AVERAGE LATENCY (MSEC) UNDER NO-HANDOVER AND HANDOVER CONDITIONS WITH AND WITHOUT IBC.

n	d_{V_i} = 50 B				d_{V_i} = 30 KB				d_{V_i} = 300 KB				d_{V_i} = 3 MB			
	No HO		HO		No HO		HO		No HO		HO		No HO		HO	
	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/	w/o	w/
10	21.0	23.2	24.1	26.8	28.8	29.0	34.6	36.3	40.3	45.6	53.4	63.7	48.2	54.5	66.2	78.4
40	72.6	83.0	85.3	99.3	69.7	77.1	90.6	101.8	110.2	133.5	148.3	190.5	120.8	130.9	172.4	188.7
70	119.5	119.1	140.8	144.1	126.7	129.6	171.0	178.9	134.0	151.0	190.0	216.2	137.0	153.6	196.5	222.8
100	216.3	229.2	255.9	277.5	234.7	243.2	328.5	366.3	257.0	272.0	370.0	394.0	260.0	278.0	378.0	407.0

To isolate the amplification factor, we define the handover amplification ratio:

$$\eta_{\text{HO}} = \frac{T_{\text{total}}^{\text{handover}}}{T_{\text{total}}^{\text{no handover}}}. \tag{4.2}$$

Across dense scenarios ($n \geq 70$), η_{HO} consistently exceeds 1.3 for medium workloads and approaches 1.46 for large workloads, confirming that mobility induces a multiplicative latency effect when authentication is enforced. Fig. 4.4 further shows that authentication frequency nearly doubles under handover in high-density settings, reinforcing that mobility not only increases delay but also elevates cryptographic load.

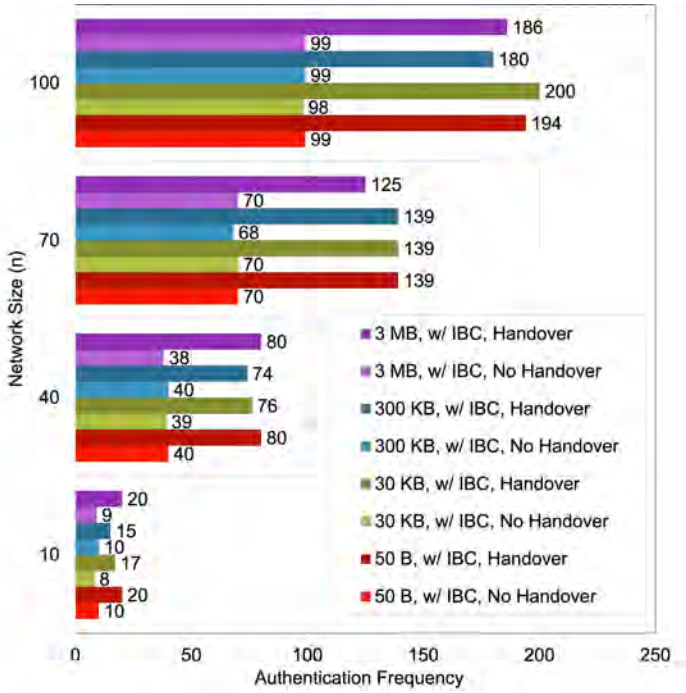

Figure 4.4. Number of authentication events under no-handover and handover conditions across network sizes and task scales.

These results demonstrate that mobility is not a secondary perturbation. It is a first-order delay amplifier once authentication is integrated. Consequently, any offloading policy that ignores handover probability will systematically underestimate execution latency and overestimate deadline satisfaction.

4.2 DT Predictive Synchronization and Deployment Optimization

Track 2 evaluates the quantitative impact of DT-enabled predictive synchronization and cyber-layer optimization across CT, ET, and HT deployments. Building upon the unified architecture introduced in Chapter 2, this section validates the central hypothesis that reducing cyber-physical misalignment enhances decision timeliness, stabilizes offloading policies, and improves scalability under density and mobility variation. The evaluation proceeds along two tightly coupled dimensions: (i) synchronization fidelity, quantified through the DT lag metric $\Delta\tau$, and (ii) deployment-level performance across CT, ET, and HT compared to traditional baselines TC, TE, and TH.

4.2.1 DT Lag Reduction and Synchronization Fidelity

DTs operate as cyber-layer state estimators. Their effectiveness depends on the temporal alignment between physical system evolution and cyber-layer representation. We quantify this alignment using the DT lag metric:

$$\Delta\tau = t_{\text{physical}} - t_{\text{DT}}, \tag{4.3}$$

where $\Delta\tau$ captures the temporal misalignment between real-world state updates and the DT state used for decision-making. Table 4.3 shows witnessed communication delay under two configurations: baseline synchronization and predictive synchronization using learning-based state estimation. The results reveal a density-sensitive lag escalation under baseline synchronization. At $n = 40$, the lag reaches 33.15 sec, reflecting congestion-induced staleness in cyber-layer state acquisition. In contrast, predictive synchronization reduces lag to 2.11 sec, corresponding to a 93.63% reduction. This reduction is not incremental. Without prediction, the cyber layer operates on stale state information, leading to delayed or suboptimal decisions. With predictive synchronization, the cyber layer maintains

Table 4.3
Witnessed Communication Delay $\Delta\tau$ (sec): when no DT is used and when DT with prediction is used

n	**No DT**	**DT and Prediction**	**Improvement(%)**
2	1.657793333	0.196686667	88.13
5	4.144483333	0.347716667	91.61
10	8.288966667	0.599433333	92.76
15	12.43345	0.85115	93.15
20	16.57793333	1.102866667	93.34
25	20.72241667	1.354583333	93.46
30	24.8669	1.6063	93.54
35	29.01138333	1.858016667	93.59
40	33.15586667	2.109733333	93.63

forward-estimated state alignment, enabling anticipatory orchestration. Beyond mean lag reduction, predictive synchronization significantly reduces lag dispersion over time. Baseline synchronization exhibits increasing variance as network density grows, while predictive DT narrows the dispersion envelope of $\Delta\tau$. Reduced variance improves decision confidence, particularly for handover anticipation and deadline-sensitive scheduling.

4.2.2 Deployment Performance Across CT, ET, and HT

Having established improved temporal alignment, we evaluate its system-level impact on offloading performance. Fig. 4.7 reports latency as a function of vehicle density across CT, ET, and HT compared to TC, TE, and TH. DT-assisted deployments consistently exhibit moderated latency growth relative to traditional baselines. CT achieves the lowest absolute latency, ranging between 16.31–19.35 msec, with improvements up to 33.4% compared to TC. ET improves latency by up to 28.6% relative to TE. HT demonstrates the strongest scalability behavior under density variation, achieving up to 37.5% improvement over TH. Traditional deployments exhibit steep latency slopes as n increases due to reactive scheduling and delayed state updates. In contrast, DT-assisted deployments benefit from predictive coordination, preventing queue buildup and avoiding execution under stale system estimates.

Task-size scaling results, shown in Fig. 4.5, further confirm this behavior. For large workloads (3.1 MB) at $n = 100$, CT reduces latency by up to 70.7% compared to TC. ET achieves up to 49.6% reduction relative to TE, while HT achieves up to 38.5% improvement over TH. Latency growth

under DT-assisted deployments is sublinear relative to traditional models, reflecting improved load distribution and anticipatory allocation.

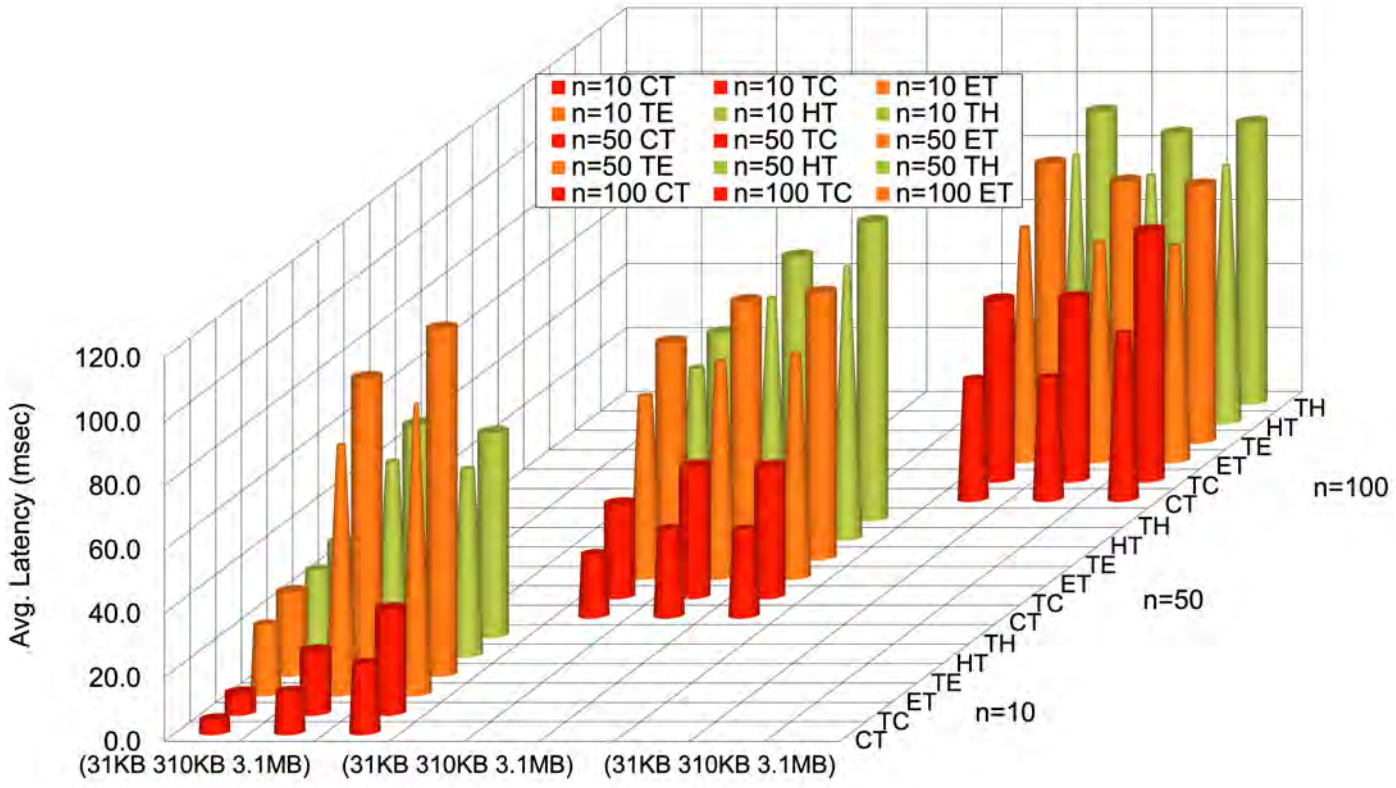

Figure 4.5. Average latency (msec) vs. task size d_{V_i} (KB) at different densities.

Transmission-rate scaling results in Fig. 4.6 demonstrate proportional latency reduction under increased data rates. When access rates increase from 200/400 Mbps to 600/1200 Mbps, CT achieves up to 47.2% latency reduction at $n = 100$. HT leverages both cloud and edge upgrades adaptively, confirming that predictive DT coordination enables effective utilization of available communication capacity.

These results establish three structural system transformations enabled by predictive DT intelligence:

- **Temporal Realignment**: Predictive synchronization reduces $\Delta\tau$ by up to 93.6%, preventing stale-state decisions.

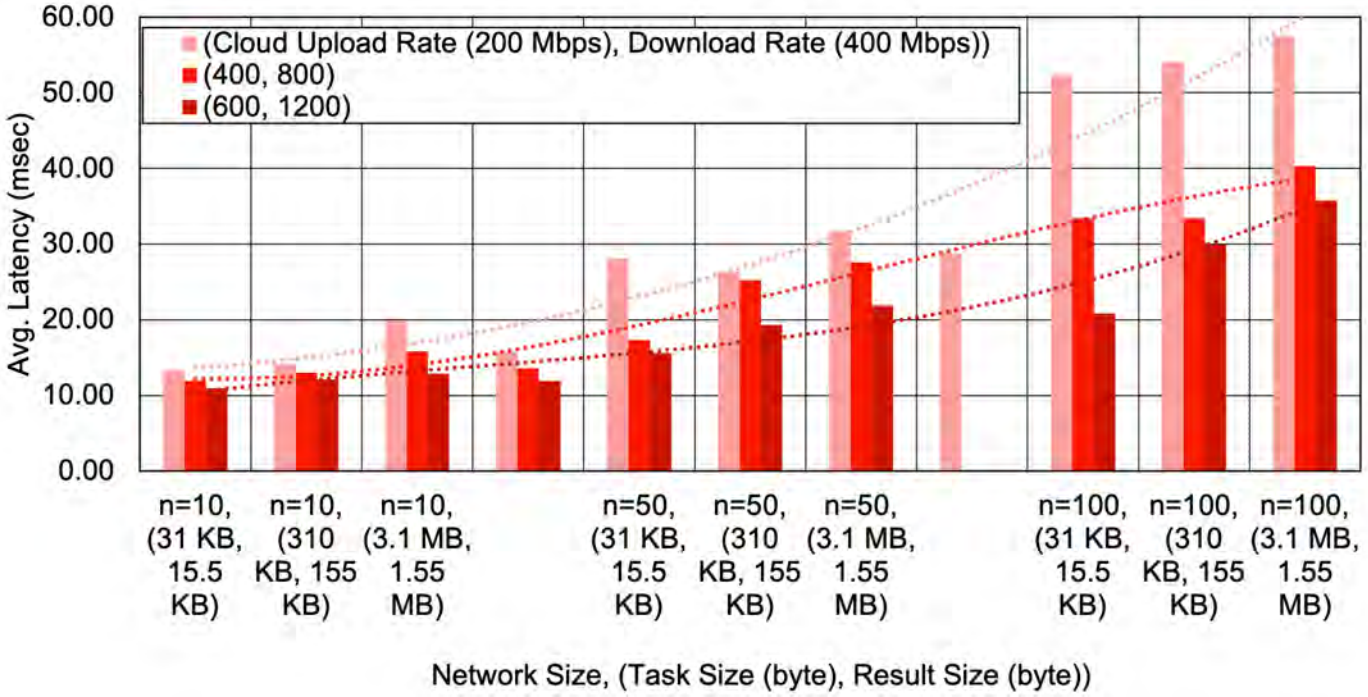

Figure 4.6. CT performance vs. data rates at various network and task sizes.

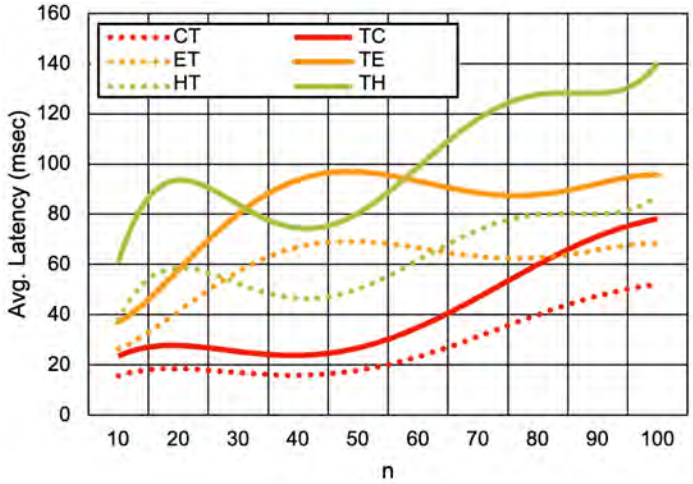

Figure 4.7. Average latency (msec) vs. number of vehicles n.

- **Decision Robustness**: Reduced lag variance stabilizes offloading policies under density and mobility variation.

- **Scalable Optimization**: CT, ET, and HT consistently outperform TC, TE, and TH across density, task size, and transmission-rate scaling.

HT emerges as the most adaptable configuration under heterogeneous and dynamic conditions, combining edge proximity with cloud capacity while leveraging predictive state alignment.

Importantly, Track 1 demonstrated that authentication overhead introduces density-sensitive latency inflation. Track 2 shows that predictive DT intelligence compensates for this inflation by reducing temporal misalignment and enabling anticipatory decision-making. Thus, DT-enabled optimization acts as a structural stabilizer in security-constrained vehicular systems.

4.3 RL-Driven Safety Under Behavioral Uncertainty

Track 3 evaluates RL-based safety control for autonomous vehicle navigation at unmarked crosswalks under stochastic pedestrian hesitation. The scenario, MDP formulation, reward shaping, and PPO-LSTM design, while metrics and reporting protocol follow Chapter 3, Section 3.5. The central claim validated here is that safety-critical interaction can be achieved under behavioral uncertainty without trajectory prediction models or multi-agent co-learning, using an AV-centered PPO-LSTM policy that learns temporal risk patterns induced by hesitation.

4.3.1 Training Convergence and Policy Stabilization

Training stability is assessed using accumulated reward trajectories across the five seeds. Fig. 4.8 shows the learning curves under ($\gamma = 0.98,$ entropy coefficient $= 0.01$) over 50,000 frames. Early-stage behavior exhibits high variance due to exploration and the pedestrian's stochastic hesitation, including pauses and reversals near the crosswalk center. After approximately 3,000 environment steps, reward trends stabilize across runs, indicating consistent policy refinement rather than oscillatory behavior.

This stabilization is attributed to three interacting design choices. First, the PPO clipped objective constrains destructive policy updates under noisy returns. Second, the LSTM memory captures temporally extended cues (e.g., repeated hesitation patterns) that are not recoverable from instantaneous observations alone. Third, entropy regularization prevents premature collapse into a single overly conservative strategy (e.g., always braking), maintaining maneuver diversity while convergence progresses.

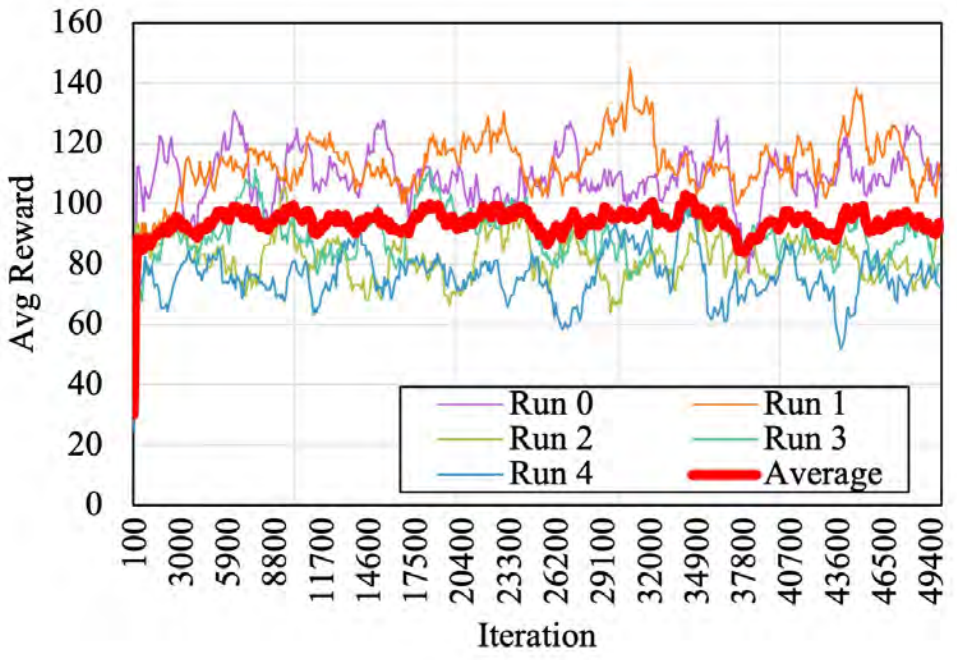

Figure 4.8. Convergence trends of average accumulated rewards over 50,000 training frames for five independent PPO-LSTM training runs under ($\gamma = 0.98$, entropy coefficient $= 0.01$).

4.3.2 Safety Robustness: Collision Rate Versus Hesitation

Safety is quantified using collision rate $\mathcal{C}_R$, defined as the ratio of collision episodes to total episodes for each ϕ_h, averaged across five seeds. Fig. 4.9a reports a monotonic increase in collision rate as hesitation grows. Under low uncertainty ($\phi_h = 0.1$), the PPO-LSTM policy achieves:

$$\mathcal{C}_R(\phi_h = 0.1) = 0.0584. \tag{4.4}$$

As the pedestrian becomes more unpredictable, collision rates increase:

$$\mathcal{C}_R(\phi_h = 0.3) = 0.1927, \qquad \mathcal{C}_R(\phi_h = 0.6) = 0.1946. \tag{4.5}$$

A key observation is the saturation effect from $\phi_h = 0.3$ to $\phi_h = 0.6$: once behavioral randomness crosses a critical level, additional hesitation does not proportionally degrade safety. This suggests that the learned policy converges to a stable avoidance regime that remains effective even when pedestrian actions are highly stochastic.

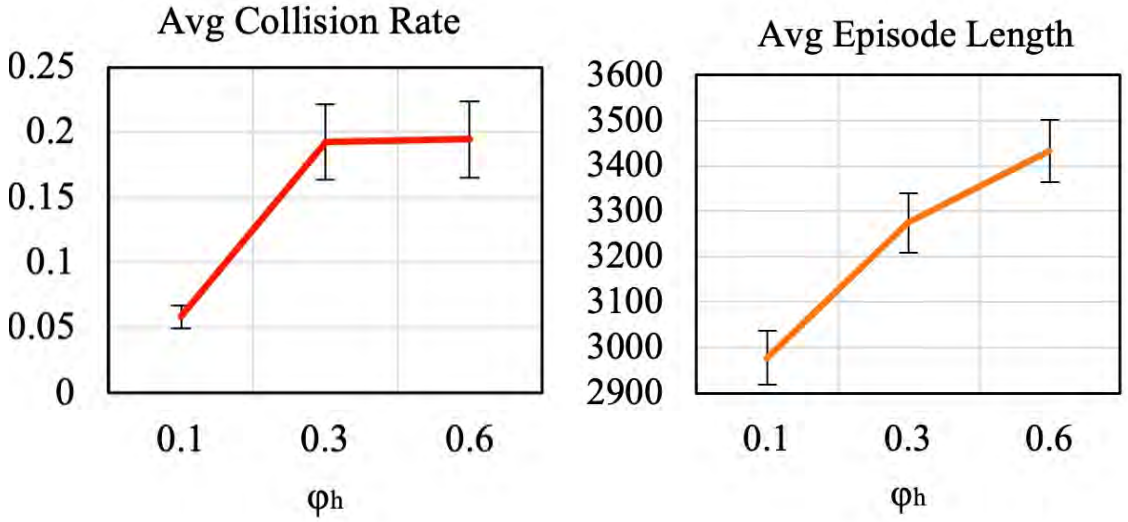

(a) Collision rate ($\mathcal{C}_R$) vs. ϕ_h.

(b) Episode length ($\mathcal{E}_L$) vs. ϕ_h.

Figure 4.9. Evaluation of safety and policy behavior for ($\gamma = 0.98$, entropy coefficient $= 0.01$) across varying pedestrian hesitation probabilities.

4.3.3 Adaptive Decision-Making: Action Distribution Under Uncertainty

Behavioral adaptation is evaluated via action distribution $\mathcal{A}_D$, which captures how the policy reallocates maneuver preferences as uncertainty changes. Fig. 4.10 shows a clear shift toward evasive maneuvers as ϕ_h increases. At $\phi_h = 0.1$, the policy already exhibits context sensitivity, with substantial use of braking and lateral dodging rather than defaulting to maintain-speed. As uncertainty increases to $\phi_h = 0.3$, the distribution becomes more uniform, reflecting strategy diversification when the pedestrian is less predictable. At $\phi_h = 0.6$, the policy increases lateral avoidance, with dodge-left emerging as the most frequent action, while maintain-speed decreases. To compactly summarize

cross-setting behavior, Table 4.4 reports:

$$\mathcal{A}_D = 22.97\% \text{ straight}, \qquad 77.03\% \text{ evasive}. \tag{4.6}$$

Table 4.4
Agent Assessment Summary

Goal	KPI	Value
Safety	Collision Rate ($\mathcal{C}_R$)	0.1476
Adaptation	Action Use ($\mathcal{A}_D$)	22.97% straight, 77.03% evasive
Efficiency	Episode Length ($\mathcal{E}_L$)	3228.67 steps

This confirms a risk-aware policy. The AV increases spatial separation through lateral maneuvers as uncertainty rises, rather than relying solely on braking. This behavior is consistent with a learned internalization of hesitation as elevated collision probability, where lateral displacement provides robustness when pedestrian motion exhibits reversals and stop-and-go dynamics.

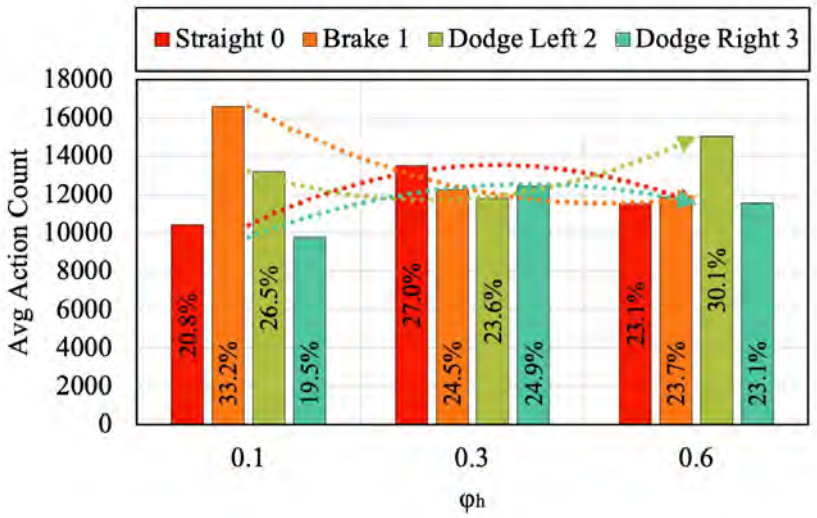

Figure 4.10. Action distribution ($\mathcal{A}_D$) across different hesitation probabilities ϕ_h. The agent prefers straight movement but exhibits increased lateral evasive actions as uncertainty grows.

4.3.4 Efficiency: Episode Length and the Safety-Efficiency Tradeoff

Efficiency is quantified via episode length $\mathcal{E}_L$, representing the total number of steps required to complete an episode. Fig. 4.9b shows that episode duration increases with hesitation. Over the last 100 episodes, total episode lengths increase from 2978 steps at $\phi_h = 0.1$ to 3275 at $\phi_h = 0.3$ and 3433 at $\phi_h = 0.6$. Relative to the baseline $\phi_h = 0.1$, this corresponds to approximately 1.099× and 1.15× longer

duration at $\phi_h = 0.3$ and $\phi_h = 0.6$, respectively. This increase reflects a standard safety-efficiency tradeoff: as pedestrian behavior becomes less predictable, the policy adopts more cautious motion (braking and evasive maneuvers), which extends episodes. Importantly, this caution is structured and adaptive, not unstable. The policy does not exhibit erratic oscillations between conflicting actions; instead, it increases conservative behavior in a consistent manner as uncertainty rises.

Track 3 validates that an AV-centered PPO-LSTM policy can learn robust safety behavior under stochastic pedestrian hesitation without trajectory prediction or multi-agent co-learning. Across increasing uncertainty ϕ_h, the policy converges reliably, maintains bounded safety degradation, and adapts action preference toward lateral evasion while accepting longer episodes as the cost of caution. Quantitatively, the agent achieves $\mathcal{C}_R = 5.84\%$ at $\phi_h = 0.1$, a cross-setting mean collision rate $\overline{\mathcal{C}_R} = 0.1476$, executes evasive actions in 77.03% of decisions, and completes episodes in an average of 3228.67 steps. Overall, the results demonstrate that memory-augmented policy optimization yields a lightweight, interpretable, and deployment-aligned safety controller for uncertain AV-pedestrian interaction, suitable for real-time constraints where explicit intent prediction or co-learning with pedestrians is impractical.

Chapter 5
Discussion and Implications

The results show that secure vehicular intelligence cannot be decomposed into isolated optimizations. Security, synchronization, and learning interact through latency, mobility, and uncertainty channels. Authentication is revealed as a structural component of end-to-end latency. Cryptographic signing and verification introduce a persistent baseline delay, while mobility amplifies this effect by converting authentication into a recurrent load factor during handovers. When authentication and mobility are excluded from latency models, delay is systematically underestimated. This means secure V2X communication must be embedded directly in scheduling and offloading optimization rather than treated as an overlay. DT synchronization operates as a predictive coordination surface. Reducing cyber-physical misalignment $\Delta\tau$ improves decision timeliness and stability across CT, ET, and HT deployments. The performance gains arise from anticipatory orchestration enabled by predictive state awareness. Synchronization fidelity therefore directly influences scalability under increasing density. Uncertainty-aware learning governs safety-critical interaction. Explicit modeling of pedestrian hesitation confirms that stochastic behavior defines the operational regime. The PPO-LSTM policy adapts conservatively as uncertainty increases, preserving bounded collision rates while sacrificing efficiency when necessary. Adaptive control under uncertainty must therefore be integrated with secure communication and synchronized state estimation to maintain robustness. Together, these findings reveal a layered dependency:

1. Authentication and mobility shape the latency structure.

2. Latency structure constrains synchronization fidelity.

3. Synchronization fidelity determines decision timeliness.

4. Decision timeliness influences safety performance under uncertainty.

5.1 Scientific Output and Broader Impact

The project produced peer-reviewed dissemination across IEEE conferences and a Q1 journal [46–51]. Recognition includes:

- **Best Runner-Up Paper Award:** SmartNets 2025 Conference Presentation, "When Pedestrians Hesitate: PPO-Based RL Collision Avoidance in Uncertain Scenarios," Istanbul, July 22–24, 2025.

- **OSU CAR News Feature:** "Evasive Maneuvers: New AI Training Program Ensures Safer Travel with Hesitant Pedestrians," September 2025.

- **Invited Seminar:** Al-Shareeda, S., "Shifting Ground: How AI Transforms Skills, Research, and the Competence Ecosystem," DevFest Basra 2025, organized by GDG Basra and Women Techmakers Basra, Basrah, Iraq, December 2025.

These outcomes demonstrate technical relevance beyond academic publication and contribute to workforce development in secure intelligent transportation systems.

5.2 Limitations

The results are derived from controlled simulation frameworks with abstractions necessary for isolating causal effects. The primary limitations are:

- **Mobility and channel modeling:** Urban propagation complexity, blockage dynamics, and full cellular scheduling stacks are abstracted.

- **Queueing and compute contention:** Resource allocation is modeled through simplified capacity abstractions rather than full virtualization and multi-tenant interference effects.

- **Authentication modeling:** Overhead is represented via cryptographic cycle and payload models rather than complete protocol-layer workflows including certificate distribution, revocation, and handshake signaling.
- **Safety environment scope:** The RL environment models structured pedestrian hesitation but does not incorporate full multi-agent interaction, perception-stack uncertainty, or occlusion dynamics.
- **Deployment readiness:** Real-world implementation would require safety shields, runtime monitors, and hardware-in-the-loop validation.

Despite these abstractions, the structural dependencies across authentication, mobility, synchronization, and uncertainty remain consistent. Coordinated cyber-layer intelligence mitigates these interactions more effectively than isolated optimization.

Chapter 6
Conclusion

This work established a unified cyber-layer intelligence framework for secure, scalable, and uncertainty-aware vehicular systems. The project addressed three tightly coupled challenges: authentication-induced latency under mobility, real-time computation offloading under dynamic network conditions, and safety-critical decision-making under behavioral uncertainty.

First, authentication is not a negligible additive cost but a structural latency component whose impact scales with density and mobility. Under C-V2X handovers, reauthentication and forwarding amplify delay and increase authentication frequency significantly. Security must therefore be mobility-aware and integrated within optimization loops. Second, DT architectures function as predictive control layers. By reducing cyber-physical misalignment ($\Delta\tau$), DT synchronization improves offloading timeliness and scalability. CT, ET, and HT deployments outperform reactive baselines when prediction is embedded into orchestration. Third, uncertainty-aware RL enables adaptive safety behavior under stochastic pedestrian hesitation. The PPO-LSTM framework maintains bounded collision rates without explicit trajectory prediction, demonstrating that learning-based control can operate under behavioral ambiguity. The central insight is that secure vehicular intelligence is a coupled cyber-physical problem. Cryptographic overhead, congestion scaling, mobility transitions, synchronization delay, and behavioral uncertainty interact nonlinearly. Coordinated cyber-layer intelligence provides a mechanism to manage these interactions coherently.

Bibliography

[1] Sarah Al-Shareeda and Fusun Ozguner. Alternating authentications to match the situational context of an intelligent communicating vehicle. *Vehicular Communications*, 23:100248, 2020.

[2] Guanjie Cheng, Junqin Huang, Yewei Wang, Jun Zhao, Linghe Kong, Shuiguang Deng, and Xueqiang Yan. Conditional privacy-preserving multi-domain authentication and pseudonym management for 6g-enabled iov. *IEEE TIFS*, pages 1–1, 2023.

[3] Pandi Vijayakumar, Maria Azees, Sergei A. Kozlov, and Joel J. P. C. Rodrigues. An anonymous batch authentication and key exchange protocols for 6g enabled vanets. *IEEE Transactions on Intelligent Transportation Systems*, 23(2):1630–1638, 2022.

[4] Homa Maleki, Mehmet Başaran, and Lütfiye Durak-Ata. Handover-enabled dynamic computation offloading for vehicular edge computing networks. *IEEE TVT*, 72(7):9394–9405, 2023.

[5] Yuzhi Zhou, Jinlong Sun, Jie Yang, Guan Gui, Haris Gacanin, and Fumiyuki Adachi. Handover strategy based on side information in air-ground integrated vehicular networks. *IEEE Transactions on Vehicular Technology*, 71(10):10823–10831, 2022.

[6] Syed Danial Ali Shah, Mark A Gregory, and Shuo Li. A distributed control plane architecture for handover management in mec-enabled vehicular networks. In *2021 31st International Telecommunication Networks and Applications Conference (ITNAC)*, pages 188–191. IEEE, 2021.

[7] Shunrong Jiang, Jinpeng Li, Guohuai Sang, Haiqin Wu, and Yong Zhou. Vehicular edge computing meets cache: An access control scheme with fair incentives for privacy-aware content delivery. *IEEE Transactions on Intelligent Transportation Systems*, 25(8):8404–8418, 2024.

[8] Juliang Cai, Xiaofeng Tao, and Chenyu Wang. Cooperative authentication scheme for heterogeneous networks based on identity group signature and blockchain. *IEEE Transactions on Vehicular Technology*, 73(1):1394–1399, 2024.

[9] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In *Advances in Cryptology–CRYPTO 2004*, pages 41–55. Springer, 2004.

[10] Sarah Al-Shareeda and Fusun Ozguner. Lightweight identity-based cryptographic framework for smart vehicles: Cryptographic proof. In *Smart Cities Symposium 2018*, pages 1–6, 2018.

[11] Sarah Al-Shareeda, Sema F. Oktug, Yusuf Yaslan, Gokhan Yurdakul, and Berk Canberk. Does twinning vehicular networks enhance their performance in dense areas? *arXiv preprint arXiv:2402.10701*, 2024.

[12] Sarah Al-Shareeda, Khayal Huseynov, Lal Verda Cakir, Craig Thomson, Mehmet Ozdem, and Berk Canberk. *AI-based traffic analysis in digital twin networks*, pages 83–132. Telecommunications. Institution of Engineering and Technology, 2024.

[13] Wenhao Fan, Yaoyin Zhang, Guangtao Zhou, and Yuan'an Liu. Deep reinforcement learning-based task offloading for vehicular edge computing with flexible rsu-rsu cooperation. *IEEE Transactions on Intelligent Transportation Systems*, 25(7):7712–7725, 2024.

[14] Xuanhong Zhou, Muhammad Bilal, Ruihan Dou, Joel J. P. C. Rodrigues, Qingzhan Zhao, Jianguo Dai, and Xiaolong Xu. Edge computation offloading with content caching in 6g-enabled iov. *IEEE Transactions on Intelligent Transportation Systems*, 25(3):2733–2747, 2024.

[15] Namory Fofana, Asma Ben Letaifa, and Abderrezak Rachedi. Intelligent task offloading in vehicular networks: a deep reinforcement learning perspective. *IEEE TVT*, pages 1–16, 2024.

[16] Rui Men, Xiumei Fan, Kok-Lim Alvin Yau, Axida Shan, and Gang Yuan. Hierarchical aerial computing for task offloading and resource allocation in 6g-enabled vehicular networks. *IEEE TNSE*, 11(4):3891–3904, 2024.

[17] Hang Shen, Yibo Tian, Tianjing Wang, and Guangwei Bai. Slicing-based task offloading in space-air-ground integrated vehicular networks. *IEEE TMC*, 23(5):4009–4024, 2024.

[18] Bin Li, Wancheng Xie, Yinghui Ye, Lei Liu, and Zesong Fei. Flexedge: Digital twin-enabled task offloading for uav-aided vehicular edge computing. *IEEE Transactions on Vehicular Technology*, 2023.

[19] Qian Liu, Rui Luo, Hairong Liang, and Qilie Liu. Energy-efficient joint computation offloading and resource allocation strategy for isac-aided 6g v2x networks. *IEEE TGCN*, 7(1):413–423, 2023.

[20] Muhammet Hevesli, Abegaz Mohammed Seid, Aiman Erbad, and Mohamed Abdallah. Task offloading optimization in digital twin assisted mec-enabled air-ground iiot 6g networks. *IEEE Transactions on Vehicular Technology*, pages 1–16, 2024.

[21] Stefan Forsström and Ye Yuhang. A testbed for evaluating task offloading algorithms in edge-fog-cloud v2i scenarios. In *19th Swedish National Computer Networking and Cloud Computing Workshop (SNCNW 2024)*, 2024.

[22] Qianpiao Ma, Hongli Xu, Haibo Wang, Yang Xu, Qingmin Jia, and Chunming Qiao. Fully distributed task offloading in vehicular edge computing. *IEEE Transactions on Vehicular Technology*, 73(4):5630–5646, 2024.

[23] Lal Verda Cakir, Sarah Al-Shareeda, Sema F Oktug, Mehmet Özdem, Matthew Broadbent, and Berk Canberk. How to synchronize digital twins? a communication performance analysis. In *The 28th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, pages 123–127. IEEE, 2023.

[24] Xiaobin Tan, Mingyang Wang, Tao Wang, Quan Zheng, Jun Wu, and Jian Yang. Adaptive task scheduling in digital twin empowered cloud-native vehicular networks. *IEEE Transactions on Vehicular Technology*, 73(6):8973–8987, 2024.

[25] Xinyi Nian, Zhuping Zhou, Yifei Cai, et al. A confidence-based approach to predicting pedestrian-vehicle conflicts at unsignalized mid-block crosswalks: Insights into pedestrian psychology and behavior. *SSRN: https://ssrn.com/abstract=5042765*, 2024.

[26] Tao Li, Chengxi Hu, NN Sze, Zhanbo Sun, Hongliang Ding, and Tiantian Chen. Modelling pedestrian-vehicle interaction behaviours at non-signalised crosswalks using a game theory approach incorporating the risk perception. *Transportmetrica A: Transport Science*, pages 1–31, 2025.

[27] Meiting Dang, Dezong Zhao, Yafei Wang, and Chongfeng Wei. Dynamic game-theoretical decision-making framework for vehicle-pedestrian interaction with human bounded rationality. *IEEE Transactions on Intelligent Transportation Systems*, 2025.

[28] Amir Hossein Kalantari. *Modelling vehicle-pedestrian interactions at unsignalised locations employing game-theoretic models*. PhD thesis, University of Leeds, 2023.

[29] Lakshmi Devi Subramanian, Elizabeth E O'Neal, Nam-Yoon Kim, Megan Noonan, Jodie M Plumert, and Joseph K Kearney. Deciding when to cross in front of an autonomous vehicle: how child and adult pedestrians respond to ehmi timing and vehicle kinematics. *Accident Analysis & Prevention*, 202:107567, 2024.

[30] David R Large, Catherine Harvey, Madeline Hallewell, Xuekun Li, and Gary Burnett. On face value: a ghost driver field study investigating interactions between pedestrians and a driverless vehicle with anthropomorphic displays. *Ergonomics*, pages 1–19, 2025.

[31] E Muhammad Saim, Sarah Al-Shareeda, Keith Redmill, and Umit Ozgiiner. Safety in connected automated vehicles in the presence of vulnerable road users. 2024.

[32] Xiaoyuan Zhao. *Assessing pedestrian decision making in the presence of automated vehicles: Mitigating risks for safer urban environment*. PhD thesis, Queensland University of Technology, 2024.

[33] Rubén Izquierdo, Javier Alonso, Ola Benderius, Miguel Ángel Sotelo, and David Fernández Llorca. Pedestrian and passenger interaction with autonomous vehicles: Field study in a crosswalk scenario. *International Journal of Human–Computer Interaction*, pages 1–19, 2024.

[34] Saki Rezwana and Nicholas Lownes. Interactions and behaviors of pedestrians with autonomous vehicles: A synthesis. *Future Transportation*, 4(3):722–745, 2024.

[35] Yueyang Wang, Aravinda Ramakrishnan Srinivasan, Yee Mun Lee, and Gustav Markkula. Modeling pedestrian crossing behavior: A reinforcement learning approach with sensory motor constraints. *arXiv preprint arXiv:2409.14522*, 2024.

[36] Hamid Taghavifar and Ardashir Mohammadzadeh. Integrating deep reinforcement learning and social-behavioral cues: A new human-centric cyber-physical approach in automated vehicle decision-making. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, page 09544070241230126, 2024.

[37] Siyang Dai, Jun Liu, and Ngai-Man Cheung. Uncertainty-aware pedestrian crossing prediction via reinforcement learning. *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.

[38] Sara Pohland, Alvin Tan, Prabal Dutta, and Claire Tomlin. Stranger danger! identifying and avoiding unpredictable pedestrians in rl-based social robot navigation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 15217–15224. IEEE, 2024.

[39] E Muhammad Saim, Sarah Al-Shareeda, Keith Redmill, Umit Ozgiiner, et al. Control of automated vehicles in vehicle-pedestrian environment. 2024.

[40] Chalumpol Trararak, Trong-Thuc Hoang, and Pham Cong-Kha. Pedestrian avoidance simulation by deep reinforcement learning using webots. In *2025 International Conference on Artificial Intelligence in Information and Communication (ICAIIC)*, pages 0734–0739. IEEE, 2025.

[41] Alexandre Brunoud, Alexandre Lombard, Nicolas Gaud, and Abdeljalil Abbas-Turki. Application of hybrid deep reinforcement learning for managing connected cars at pedestrian crossings: Challenges and research directions. *Future Transportation*, 4(2):579–590, 2024.

[42] Mahsa Golchoubian, Moojan Ghafurian, Kerstin Dautenhahn, and Nasser Lashgarian Azad. Uncertainty-aware drl for autonomous vehicle crowd navigation in shared space. *IEEE Transactions on Intelligent Vehicles*, 2024.

[43] Haochong Chen, Xincheng Cao, Levent Guvenc, and Bilin Aksun-Guvenc. Deep-reinforcement-learning-based collision avoidance of autonomous driving system for vulnerable road user safety. *Electronics*, 13(10):1952, 2024.

[44] Raphael Trumpp, Harald Bayerlein, and David Gesbert. Modeling interactions of autonomous vehicles and pedestrians with deep multi-agent reinforcement learning for collision avoidance. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, pages 331–336. IEEE, 2022.

[45] Hamid Taghavifar, Chuan Hu, Chongfeng Wei, Ardashir Mohammadzadeh, and Chunwei Zhang. Behaviorally-aware multi-agent rl with dynamic optimization for autonomous driving. *IEEE Transactions on Automation Science and Engineering*, 2025.

[46] Sarah Al-Shareeda, Fusun Ozguner, and Berk Canberk. Group-signature authentication to secure task offloading in vehicular edge twin networks. In *GLOBECOM 2024 Workshop - BlockSecSDN*, page ?, 2025.

[47] Sarah Al–Shareeda, Fusun Ozguner, Keith Redmill, Trung Q. Duong, and Berk Canberk. Lightweight authenticated task offloading in 6g-cloud vehicular twin networks. In *2025 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 01–06, 2025.

[48] Sarah Al-Shareeda, Vignesh Srinivasan, Mohammad AlMudhaf, Yasser Bin Salamah, Bander Jabr, and Fusun Ozguner. Evaluating handover impact on ibc-authenticated task offloading in c-v2x vtens. In *2025 IEEE Vehicular Networking Conference (VNC)*, pages 1–8, 2025.

[49] Sarah Al-Shareeda, Yasar Celik, Bilge Bilgili, Ahmed Al-Dubai, and Berk Canberk. Accurate ai-driven emergency vehicle location tracking in healthcare its's digital twin. In *2025 5th IEEE Middle East and North Africa Communications Conference (MENACOMM)*, pages 1–8, 2025.

[50] Sarah Al-Shareeda, Nasir Saeed, Keith Redmill, Bander A. Jabr, Yasser Bin Salamah, Ahmed Al-Dubai, and Fusun Ozguner. Novel dt-assisted vehicular task offloading for cloud, edge, and hybrid deployments. *IEEE Transactions on Vehicular Technology*, pages 1–16, 2025.

[51] Sarah Al-Shareeda, Muhammad Saim, Bander Jabr, Yasser Bin Salamah, Faisal Alanazi, Gokhan Yurdakul, Fusun Ozguner, and Umit Ozguner. When pedestrians hesitate: Ppo-based rl collision avoidance in uncertain scenarios. In *2025 International Conference on Smart Applications, Communications and Networking (SmartNets)*, pages 1–7, 2025.