Sensors & Transducers, Vol. 270, Issue 3, November 2025, pp. 48-59

www.sensorsportal.com

Sensors & Transducers

Published by IFSA Publishing, S. L., 2025
http://www.sensorsportal.com

AI-enhanced VLC Cyber-physical Architecture for Adaptive Airport Traffic Management

[1, 2, 3] Manuela VIEIRA, [1, 2] Manuel A. VIEIRA, [1] Gonçalo GALVÃO, [1, 2] Paula LOURO, [1, 4] Pedro VIEIRA and [1, 2] Alessandro FANTONI

[1] Electronics Telecommunication and Computer Dept. ISEL/IPL,
R. Conselheiro Emídio Navarro, 1959-007 Lisboa, Portugal
[2] UNINOVA –CTS and LASI, Quinta da Torre, Monte da Caparica, 2829-516, Caparica, Portugal
[3] NOVA School of Science and Technology, Quinta da Torre,
Monte da Caparica, 2829-516, Caparica, Portugal
[4] Instituto de Telecomunicações, Instituto Superior Técnico, 1049-001, Lisboa, Portugal
E-mail: mv@isel.ipl.pt

Received: 29 May 2025 / Revised: 8 Nov. 2025 / Accepted: 10 Nov. 2025 / Published: 28 Nov. 2025

Abstract: Modern airports are complex Cyber-Physical Systems (CPS), where effective coordination between physical entities – like pedestrians and Autonomous Guided Vehicles (AGVs) – and computational components is essential for safety and efficiency. This study introduces a novel CPS architecture that integrates Artificial Intelligence (AI) and Visible Light Communication (VLC) to optimize mobility and enhance real-time responsiveness. Using tetrachromatic LED luminaires modulated via On-Off Keying (OOK) and amorphous SiC optical receivers in a mesh-based hybrid topology, the system creates a VLC infrastructure that delivers real-time, location-aware navigation. A custom protocol ensures low-latency, reliable data exchange between agents and the digital core. VLC receivers capture continuous data on agent positions and movements, which is processed by Deep Reinforcement Learning (DRL) agents trained via Q-learning. These agents adaptively manage traffic flow, minimize congestion, and improve throughput. Simulations and experiments confirm the system's advantages over traditional methods, enabling GPS-independent indoor navigation, efficient mixed traffic coordination, and scalable deployment within smart airport environments.

Keywords: Cyber-physical systems (CPS), Visible light communication (VLC), Internet of things (IoT), Deep reinforcement learning (DRL), Autonomous guided vehicles (AGVs).

1. Introduction

Navigating complex indoor environments such as airports, hospitals, and shopping malls presents significant challenges due to dense traffic, unfamiliar layouts, and the limitations of traditional tools like GPS, which are often ineffective indoors. These settings demand intelligent systems capable of real-time sensing, communication, and decision-making, defining characteristics of Cyber-Physical Systems (CPS).

In this context, VLC-based CPS architectures offer a promising solution by leveraging existing LED lighting infrastructure to provide both illumination and high-precision data transmission [1, 2]. VLC uses modulated visible light, imperceptible to the human eye, which is detected by optical receivers to enable accurate indoor positioning and navigation.

This work proposes a CPS framework that integrates Artificial Intelligence (AI) and Visible Light Communication (VLC) to manage airport traffic efficiently, enhancing mobility and safety for both

pedestrians and Autonomous Guided Vehicles (AGVs). The system features VLC-enabled luminaires with On-Off Keying (OOK) modulation and tandem amorphous SiC optical receivers, organized in a hybrid mesh topology to ensure coverage and data redundancy [3, 4].

Deep Reinforcement Learning (DRL) is incorporated to dynamically adapt positioning and routing strategies based on real-time environmental conditions [5-7]. DRL models learn from user interactions, obstacles, and contextual changes, enabling seamless navigation in crowded, unpredictable spaces. This adaptability supports precise guidance to gates, check-ins, and amenities, improving user experience and operational efficiency.

This work introduces an adaptive reward mechanism that refines agent behavior. It improves learning by balancing vehicle storage, congestion control, and pedestrian safety across the network [8, 9].

For airport operators, the proposed CPS ensures efficient asset tracking, streamlines operations, and supports real-time decision-making. The VLC-DRL integration, reinforced by adaptive reward modeling, improves scalability and performance, offering intelligent coordination and personalized services in complex indoor environments.

The remainder of this article is organized as follows. Section 2, presents the system design and methodological foundations, including the principles of Visible Light Communication (VLC) and the overall system architecture. Section 3, introduces the proposed VLC-DRL airport model, detailing the airport infrastructure, the DRL-based smart traffic signal control approach, and the operational dynamics of coordinated agents. Section 4, describes the AI optimization model and the rerouting strategy, including the training and testing of rerouting techniques and a comparative analysis of standard versus rerouting-enhanced scenarios. Section 5, focuses on the adaptive reward design for horizontal and vertical shared arteries, outlining the simulation setup, training performance, and analyses of AGV and pedestrian halting behaviors. Finally, Section 6, summarizes the main conclusions and outlines directions for future research.

2. System Design and Methodological Approach

This work aims to develop a VLC-based guidance system tailored for mobile users in large, multi-terminal airports.

2.1. Fundamentals of Visible Light Communication (VLC)

The system, draft in Fig. 1 consists of two core components: a transmitter and a receiver. The transmitter module uses white tetra-chromatic LEDs (WLEDs), where only one chip per luminaire transmits data using specific wavelengths – Red (626 nm), Green (530 nm), Blue (470 nm), and Violet (400 nm) – while the others maintain ambient lighting. Data is encoded using On-Off Keying (OOK) modulation, which transmits binary information via the presence or absence of light pulses. This method offers robust, low-cost, and adaptive communication that integrates easily with standard lighting infrastructure.

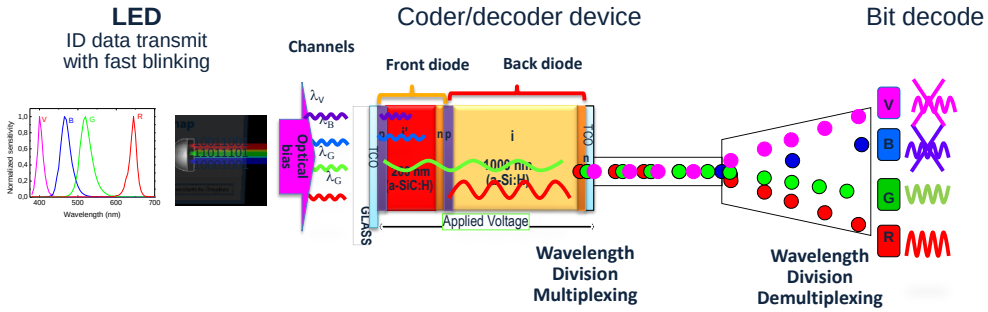

Fig. 1. Structure and operating principle of the PIN/PIN receiver.

The receiver module employs a photodetector based on a double PIN a-SiC/a-Si heterostructure, acting as an active optical filter. Designed to selectively detect wavelengths from the visible spectrum, this structure converts modulated light signals into electrical signals, which are then decoded. The sensor achieves high sensitivity, precision, and reliability in dynamic indoor environments [10].

The system operates by transmitting pulsed signals through red, green, blue, and violet light channels, each encoding distinct bit sequences. These signals are absorbed according to wavelength and analyzed via photocurrent measurements under −8 V bias and 390 nm background illumination. The receiver acts as an active optical filter, amplifying longer wavelengths (green/red) and attenuating shorter ones (blue/violet), with spectral gain exceeding unity above 500 nm.

The signal undergoes adaptive filtering, amplification, and demultiplexing to reconstruct and decode the transmitted data. For accurate positioning,

the receiver must reside within overlapping transmitter regions, generating a multiplexed (MUX) signal. Grid calibration avoids overlap conflicts, defining unique fingerprint zones for localization. The final decoded message is then relayed to the user [11].

2.2. System Architecture and Components

In VLC-based tracking, indoor spaces are divided into positioning cells managed by a Central Manager (CM), which computes optimal routes based on building layout and luminaire placement. Users request guidance and receive real-time updates via VLC signals displayed on their receivers.

A mesh cellular hybrid structure displayed in Fig. 2 underpins the system architecture, enabling secure device-to-device (D2D) communication. Each WLED luminaire emits a unique VLC signal serving as a location beacon (L2D), allowing optical receivers to identify user trajectories through a specialized positioning algorithm.

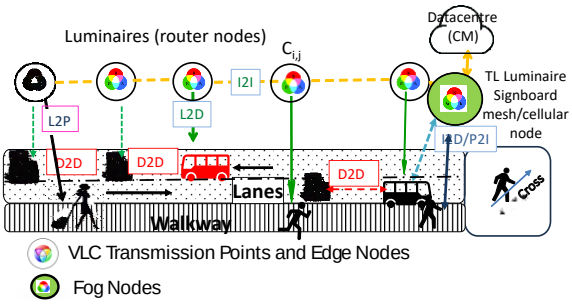

Fig. 2. Draft design of a single-lane configuration in the Edge/Frog hybrid architecture.

The computed indoor route, q(x, y, z, δ, t), captures spatial and temporal coordinates for pedestrians and AGVs, adjusted according to user type and assigned average speed. Navigation data is transmitted to users (I2D/P) via emitters embedded in smart infrastructure, such as light controllers or digital displays, enabling dynamic and context-aware wayfinding.

3. VLC-DRL Airport Model

Airport capacity is shaped by the layout of gateways, boarding zones, and door areas. Efficient allocation reduces pedestrian travel distances during arrivals, transfers, or shopping, improving service quality.

3.1. Airport Infrastructure Modeling

Key evaluation factors include walking distances and access to transport. Although pedestrian speeds in terminals are under-documented, studies indicate slower movement at junctions, entrances, or areas with signage and visual cues. These behavioral patterns must be considered in layout design.

Crowded conditions during peak hours or delays can hinder flow and safety. To address this, the system integrates adaptive rewards that adjust traffic priorities based on real-time congestion and pedestrian density. When clusters form, pedestrian phases are favored to ensure smoother, safer navigation.

Fig. 3a illustrates a simulated airport model with RGBV luminaire-based positioning (X_{ij}), pedestrian paths, and terminal layout for navigation assistance.

The simulation covers one horizontal artery and three terminals: T0 (international), T1 and T2 (domestic), each with four arms and two lanes plus sidewalks. Terminal's configuration includes lane codes (L0–7) and traffic signals (TL0–15). Traffic flows along four cardinal points, using binary request-and-response segments for decisions (e.g., turn left/straight or right). Passenger lanes, waiting zones, and crosswalks are integrated. Right lanes handle right turns and straight routes, while left lanes serve left turns. Centralized traffic lights (TL 0–15) (Fig. 3b) are controlled by a Central Manager (CM) to prevent collisions.

Data transmission in the airport VLC system uses a synchronous 64-bit frame structure with OOK modulation, where binary data is represented by turning the light on and off. Each luminaire includes RGBV WLEDs, enabling four simultaneous signals, requiring a four-channel filtering receiver.

Wavelength-calibrated signal amplitudes produce 16 optical combinations, resulting in distinct photocurrent levels weighted by photodetector gain. The PIN-PIN demultiplexer decodes the message using photon penetration depths.

Each output is mapped to a binary code based on optical gain per channel. The system supports up to 100 kbps across four channels. The decoding algorithm assigns a unique 4-bit code to the 2^4 on/off combinations [X_1, X_2, X_3, X_4], where X = 1 if the channel is active.

Data is extracted using a pre-calibrated response curve from the photodetector, followed by a parity check to ensure integrity. Decoded MUX signals and the frame structure reveal information such as pose, transmitter type, and traffic message.

Using Figs. 2 and 3, along with the communication protocol detailed and the technique for decoding calibrated signals emitted by transmitters [10].

Fig. 4 presents the MUX signals received by the receivers and the decoded optical signals. On the top of each figure, the analyzed environment is displayed to assist visual interpretation.

Fig. 4a illustrates the movement of two devices (D0 and D1) that depart from International Terminal T0 after a flight landing and head toward Domestic Terminal T2 for baggage drop-off. At T0 (Fig. 5a, top insert), the AGVs wait for the left-turn phase while the required VLC protocols are established.

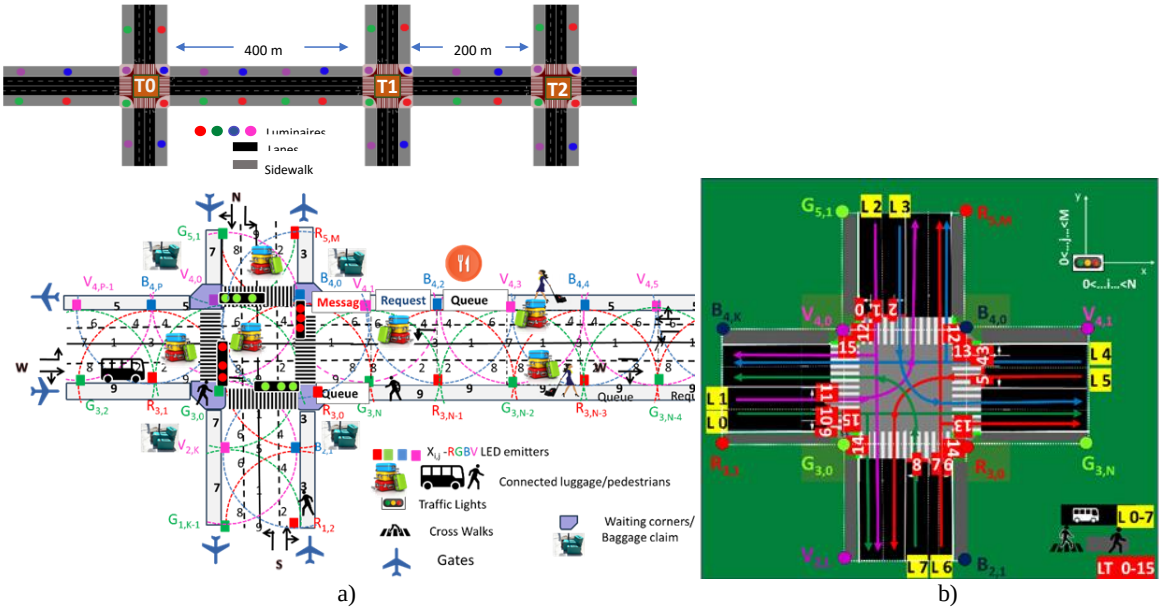

Fig. 3. a) Simulated scenario and environment, showing the optical infrastructure (Xij), generated footprints (1-9), and the flow of connected luggage and pedestrians.. b) Schematic diagram of Terminal 2, with coded lanes (L/0-7) and traffic lights (TL/0-15).

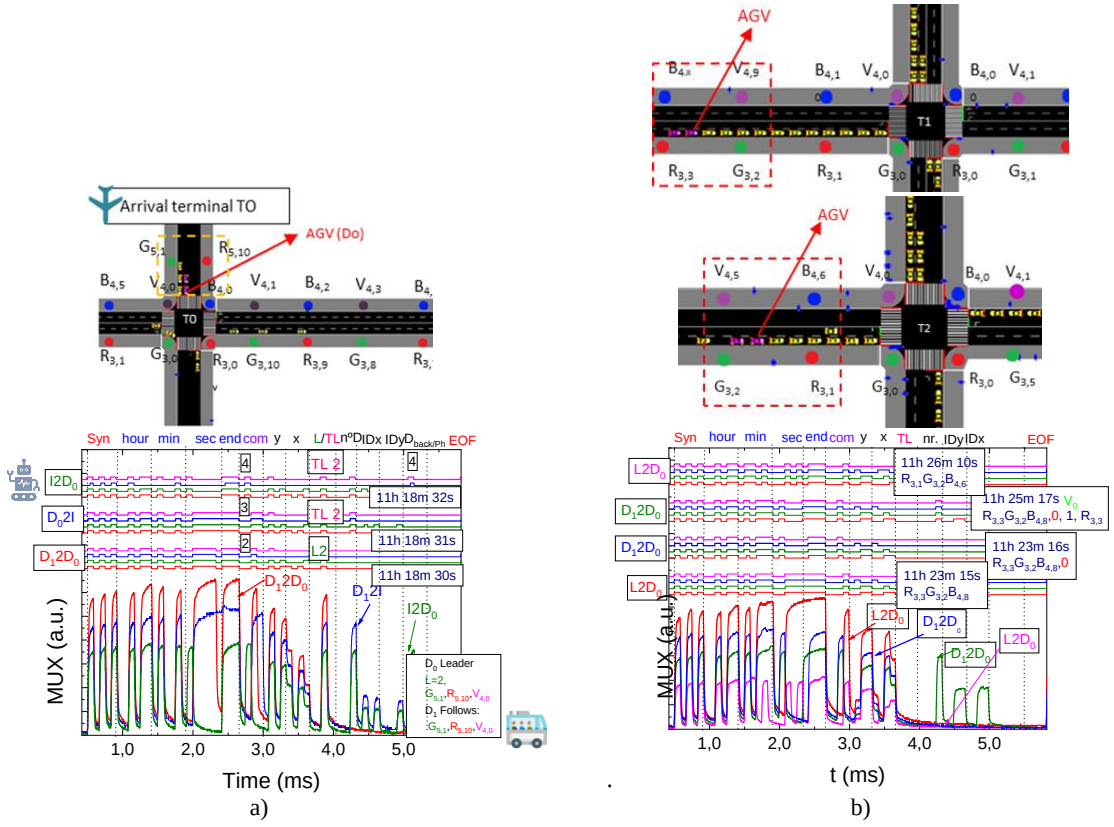

Fig. 4. MUX signal distribution across different VLC communication types.
The decoded messages are shown at the top.

In D2D communication, the rear device (D1) informs the front device (D_0) of its position ($R_{5,10}$; $G_{5,1}$; $V_{4,0}$), lane (2), follower count (0), and timestamp (11:18:30). After receiving the message, the leader D0 sends a request to the CM via $D_0$2I, reporting its own position, the traffic light ID (2), one follower, timestamp (11:18:31), follower ID, and follower queue length. The CM replies ($I2D_0$) at 11:18:32, indicating the active phase (4 – WE).

When the NS-Left phase becomes active, the devices move to Central Terminal T1 (Fig. 4b) and join the queue, receiving $L2D_0$ and D12D0 messages.

L2D0 provides lane positions ($R_{3,3}$; G3,2; B4,8) and timestamp (11:23:15). The last device transmits D12D0 information – its position, lane (0), no followers, and timestamp (11:23:16). The next device receives D2D data with its position, lane (0), one follower, timestamp (11:23:17), and the follower's position.

After T1, the devices proceed to Terminal T2, where they queue again. L2D0 and D2D communications resume, providing updated lane positions (e.g., $R_{3,1}$; $G_{3,2}$; $B_{4,6}$) and timestamps (e.g., 11:26:10).

Fig. 5 illustrates a scenario in which a flight has just arrived at International Terminal T0. Devices D0 and D1 depart from this terminal toward Domestic Terminal T2, while pedestrians at T0 wait for their phase to perform a left turn. During this period, several VLC communication procedures are triggered.

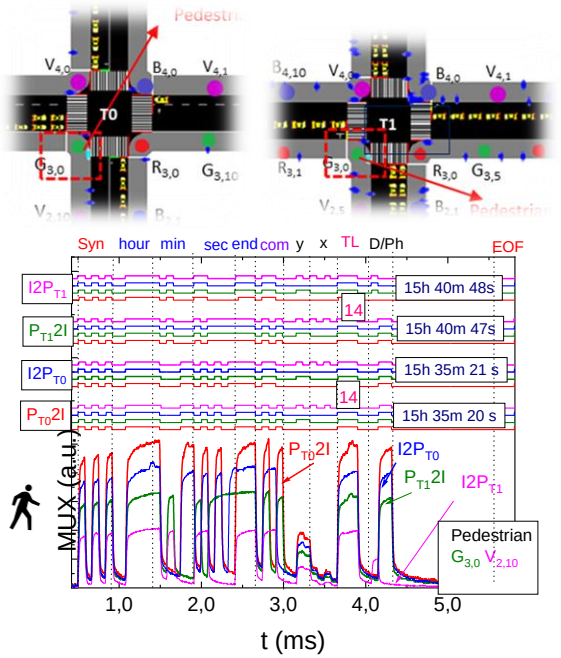

Fig. 5. MUX signal distribution across different VLC communication types. The decoded messages are shown at the top.

In the P2I exchange, the pedestrian sends their position ($G_{3,0}$; $V_{2,10}$), the traffic light ID (14), the intended direction (East), and the timestamp (15:35:20). The IM replies (I2P) with the pedestrian's ID, the same traffic light, the active phase (NS Left), and a timestamp (15:35:21). After crossing T0, the pedestrian reaches intersection T1, where another P2I/I2P sequence occurs. The P2I message again includes the position, traffic light, direction, and timestamp (15:40:47), and the IM responds with the active WE phase and the corresponding timestamp (15:40:48). Upon reaching T2, the pedestrian may later initiate a new crossing request, generating the final P2I and I2P messages.

These results demonstrate that VLC effectively supports detailed Vehicle-to-Infrastructure, Vehicle-to-Vehicle, Pedestrian-to-Infrastructure, and Infrastructure-to-Pedestrian exchanges across multiple intersections, enabling a structured and coordinated multimodal traffic management framework.

3.2. Smart Traffic Signal Control Based on Deep Reinforcement Learning

A VLC-based intelligent system enables real-time traffic management for AGVs and pedestrians using Multi-Agent Reinforcement Learning (MARL). Each terminal agent (T0, T1, T2) collects environmental data and shares traffic experiences to train a unified Deep-Q Network (DQN), which predicts optimal signal phases by minimizing waiting times.

Rewards are based on the accumulated time vehicles and pedestrians remain nearly stationary, with Eq. (1) quantifying this metric.

$$atwt_{veh,t} = \sum_{veh=1}^{n} wt_{(veh,t)}, \quad atwt_{ped,t} = \sum_{ped=1}^{n} wt_{(ped,t)} \tag{1}$$

The final reward (Eq. (2)) reflects the difference in total waiting times between time steps, weighted by priority factors (p_{veh} and p_{ped}) to allow flexible control strategies.

$$r_t = p_{veh}\left(atwt_{veh,t-1} - atwt_{veh,t}\right) + p_{ped}\left(atwt_{ped,t-1} - atwt_{ped,t}\right) \tag{2}$$

Agents store interactions in replay memory to improve training stability. Centralized methods use a single global agent, while decentralized approaches train each intersection agent independently, adapting to local traffic dynamics. The flowchart during simulation and training is displayed in Fig. 6.

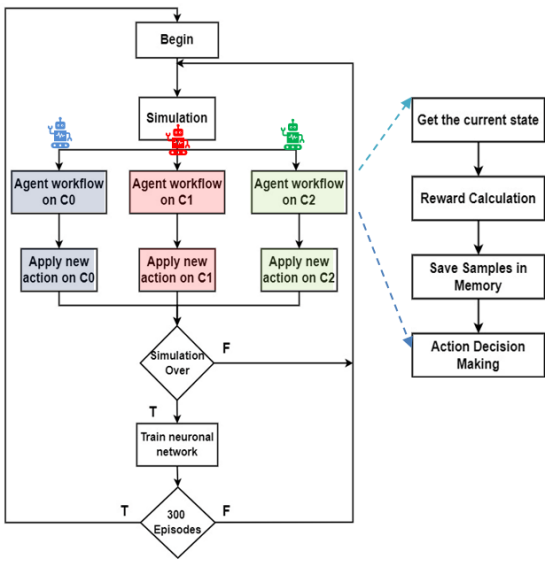

Fig. 6. Flowchart during simulation and training.

3.3. Operational Agent Dynamics in Coordinated Traffic Systems

The proposed MARL system exemplifies the coordination capabilities of cyber-physical systems in airport traffic management by enabling real-time, adaptive interactions across agents.

- *I2I Coordination (Infrastructure-to-Infrastructure):* Traffic signal agents at adjacent intersections share queue and flow data, adjusting signal timings proactively to prevent bottlenecks. This cyber-physical synchronization improves AGV continuity and reduces intersection conflicts.
- *I2D/D2I Coordination (Infrastructure-to-Device and vice versa):* Traffic agents interact with AGVs, forecasting pedestrian densities or mission urgencies. For example, an agent can reroute AGVs during pedestrian surges, or grant early green phases based on cargo priority – demonstrating real-time physical actuation informed by digital context.
- *I2P/P2I Coordination (Infrastructure-to-Pedestrian and vice versa***)**: Pedestrian agents share crossing intentions; infrastructure agents respond with adaptive signal scheduling. This ensures safety and efficiency, adjusting to dynamic human presence through seamless physical-digital integration.
- *Multi-Agent Cooperation:* In dynamic conditions (e.g., peak arrivals), agents form coalitions, sharing state information and co-adapting policies to maintain flow. This distributed CPS behavior supports resilience and scalability through collaborative control.

These CPS-driven interactions illustrate how MARL enables decentralized yet coordinated decision-making, enhancing safety, adaptability, and flow efficiency in complex airport environments.

The simulation replicates realistic airport traffic, featuring three connected four-arm intersections with two-lane bi-directional roads (Fig. 3). Each intersection is managed by an agent using a grid-based model (Fig. 7) to monitor vehicles and pedestrians across 164 cells in three layers:

- *Layer 1:* Binary cells for vehicle presence;
- *Layer 2:* Corresponding cells for normalized speed;
- *Layer 3:* Pedestrian waiting zones.

This structure offers agents a detailed spatial and dynamic understanding, aligned with VLC-based detection of vehicles and pedestrians.

Simulations were conducted using SUMO [12], an open-source microscopic traffic platform, enabling precise modeling of vehicle and pedestrian behavior, network layouts, and traffic densities. SUMO's API integrates with external programs, allowing agents to interact with real-time traffic data and refine their control policies.

Pedestrians walk at ~1 m/s, reflecting real conditions and informing the SUMO incentive structure.

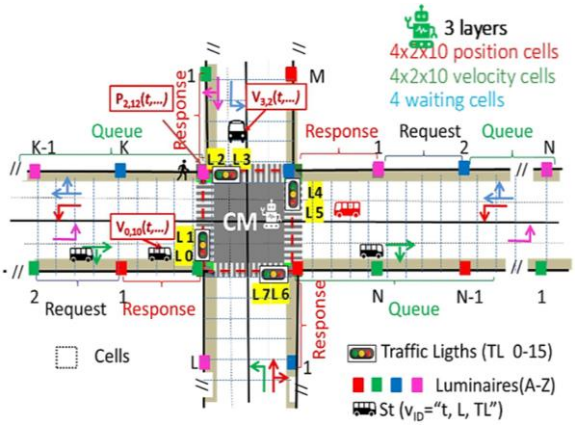

Fig. 7. State representation for each junction.

4. AI Optimization Model and Rerouting Strategy

Traffic dynamics were analyzed in single-, two-, and three-junction setups. Given the identical structure of intersections, a single neural network was trained using data from all local agents, acting as a global controller. While this centralized model effectively handles vehicle and pedestrian flows, queue peaks may occur due to the agent's limited view of neighboring queues.

To overcome this, queue length thresholds were introduced, regulating flows based on intersection capacity [13].

Three traffic scenarios were tested:

1. *Standard* – 75 % straight, 25 % turns;
2. *Symmetric Rerouting* – Redirects 75 % of vehicles via turns under high load to balance traffic;
3. *Asymmetric Rerouting* – Prioritizes flow in critical directions when one junction impacts others.

The system starts with the standard mode and switches strategies dynamically based on congestion. Although training is centralized, local agents implement rerouting strategies independently, ensuring efficient traffic regulation. Threshold-based control, informed by earlier scenario analysis, compensates for limited inter-agent communication and helps manage flow in high-density areas.

4.1. Rerouting Techniques: Network Training and Testing

To compare the effectiveness of different rerouting strategies, three neural networks were independently trained using a reward function with equal weighting, 50 % for autonomous vehicles and 50 % for pedestrians. This balanced approach ensures that both mobility aspects are considered during the learning process**.** Cumulative negative reward curves over

training episodes, comparing the performance of the three rerouting strategies are displayed in Fig. 8. The simulation involved 2600 vehicles and 2000 pedestrians over 300 episodes of 3600 seconds each [14]. Agent performance was evaluated using cumulative negative rewards, showing overall improvement despite some fluctuations due to dynamic traffic conditions.

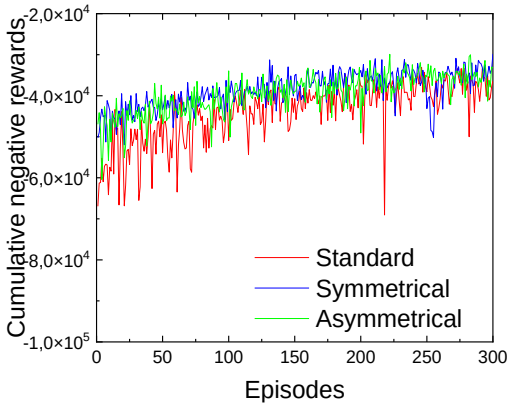

Fig. 8. Cumulative negative reward.

Both rerouting scenarios (symmetric and asymmetric) outperformed the standard case by reducing congestion and queue lengths through optimized routing. These approaches showed faster convergence, greater stability, and more consistent reward patterns. Intersection T1, located between T0 (400 m away) and T2 (200 m), emerged as the most critical due to its central position, requiring priority management to prevent operational delays and increased waiting times for both vehicles and pedestrians.

Fig. 9 show the number of stopped vehicles at T1 and T2 terminals across the three scenarios. The standard case consistently has the highest vehicle halting, while the symmetric and asymmetric rerouting strategies use advanced control to redirect vehicles and prevent terminal overload. When the vehicle thresholds – 25 between T0–T1 and 10 between T1–T2 – are reached, micro-control strategies like entry-blocking signal phases or pedestrian phases are triggered to manage congestion.

Terminal T1 consistently shows the highest vehicle accumulation, reinforcing its role as the critical hub between T0 and T2.

Fig. 10 depict pedestrian wait times, revealing peaks followed by sharp drops, indicating pedestrian phase activation. These peaks often result from either high pedestrian demand or the system prioritizing vehicles under heavy traffic.

The networks learn to optimize pedestrian phases by balancing safety with traffic flow, ensuring minimal disruption. Overall, rerouting scenarios effectively maintain both vehicle throughput and pedestrian safety.

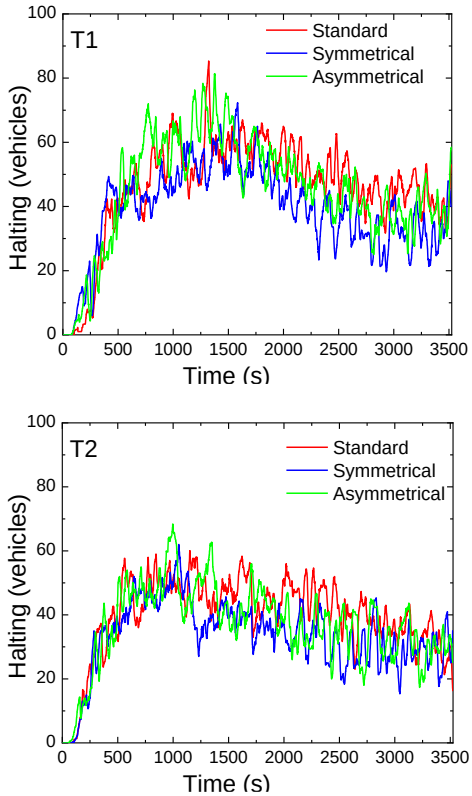

Fig. 9. Comparison of vehicle halting trends over time at intersections under standard, symmetric rerouting, and asymmetric rerouting scenarios: T1 and T2.

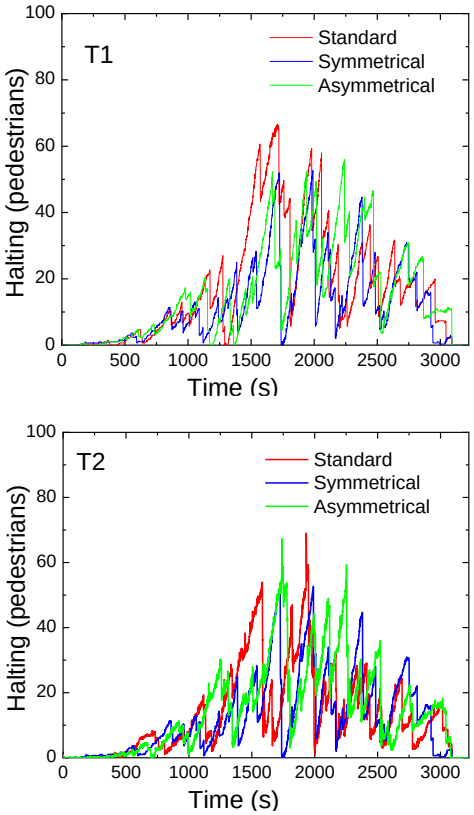

Fig. 10. Comparison of vehicle halting trends over time at intersections under standard, symmetric rerouting, and asymmetric rerouting scenarios: T1 and T2.

4.2. Comparative Analysis of Standard and Rerouting Scenarios: Activations and Traffic Management Strategies

This section evaluates the differences between the *standard* and *symmetric rerouting* scenarios in terms of phase activations and traffic management strategies.

Fig. 11 visualizes a temporal comparison of green time trends across all active phases at intersections T0, T1, and T2. a) Standard scenario. b) Symmetric rerouting scenario. c) Asymmetric rerouting scenario. Active phases are shown as an insert at the top for clarity. Results show that.

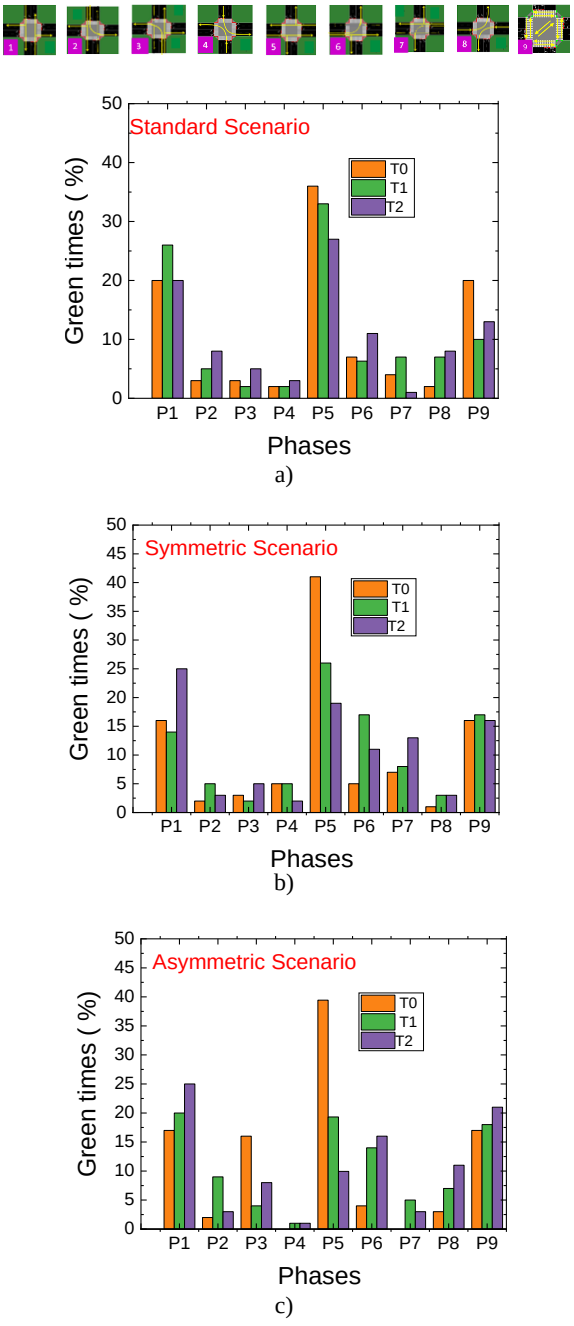

Fig. 11. Temporal comparison of green time trends across all active phases at intersections T0, T1, and T2.

Phase 1 (N-S): In the standard scenario, activated ~20 %, especially at T1 to avoid overload at T2. In rerouting scenarios, activation drops 5 % at T0 and T1 but rises at T2 to balance the high-demand 200-meter stretch.

Phase 5 (W-E): Standard scenario sees 30 % activation, mainly at T0 due to its 400-meter stretch. Activation drops at T1 and T2 to manage congestion. In symmetric rerouting, T1 and T2 see a 20 % reduction, enforcing a 15-vehicle limit. In asymmetric rerouting, T0 rises to 40 %, while T1 and T2 drop to 20 % and 10 %, focusing on E>W flow.

Phase P9 (Pedestrian): Standard scenario activates pedestrian phases 10–20 %, highest at T0. Rerouting scenarios shift activation to T1 and T2, improving vehicle flow and pedestrian management.

Phase 6 (200-Meter Clearance): Standard activation at T2 is 13 %. Rerouting increases activation slightly at T1 and T2 to reduce queues.

Phase 7 (Opposite Direction): In symmetric rerouting, activation increases to support balanced traffic across intersections.

Overall, rerouting and micro-control strategies reduce congestion, optimize pedestrian flow, and improve control of shorter roads, proving effective in managing airport traffic dynamics.

5. Adaptive Reward Design for Horizontal/Vertical Arteries

In this section, the environment is expanded to five intersections, adding complexity and realism.

5.1. Simulation Setup

As shown in Fig. 12, it now features two main arteries – one horizontal and one vertical – connected through shared intersections. The horizontal artery links W–C1–E, and the vertical artery connects N–C1–S. Intersection C1 becomes the central node, coordinating flows between both arteries and serving as the key control point in the system.

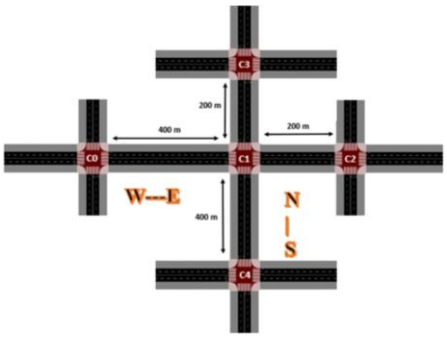

Fig. 12. Traffic scenario consisting of 5 homogeneous intersections with 4 arms each and two different arteries.

To evaluate the proposed system, an airport scenario comprising five interconnected intersections

was simulated. The environment includes 800 AGVs per hour traveling at 9 km/h and 2000 pedestrians moving at 3 km/h. Three distinct traffic generation strategies were considered:

1. *Standard strategy*, with equal AGV volumes along the W–E and N–S arteries;
2. *Circular-priority strategy*, with increased vehicle flow in the circular artery;
3. *Radial-priority strategy*, prioritizing the radial artery with higher AGV volume.

An independent DRL network was trained for each strategy, enabling a consistent comparison of learning behavior and performance.

5.2. Training Performance Analysis

To evaluate how traffic is controlled under these conditions, three neural networks were trained, one for each traffic generation strategy, and their training and testing results were analyzed.

Fig. 13 shows the training reward evolution for the standard and circular-priority strategies in the N>S artery. In both cases, the networks progressively improved their phase-selection decisions, exhibiting similar learning behavior. When comparing the standard and circular-priority strategies, clearer differences emerge, with the circular-priority model consistently achieving fewer negative rewards – indicating that environments with higher radial demand are easier for the agent to manage.

5.3. AGV Halting Behavior

Testing results in Fig. 14 illustrate AGV halting across intersections in the N>S artery, circular priority and in Fig. 15 with radial priority. Under the standard strategy, higher halting levels occur at C3 and C4 (radial artery), while the circular-priority strategy increases halting at C0 and C2 (circular artery), consistent with the corresponding traffic loads.

Regarding vehicle halting at the intersections, the most significant difference occurs at intersection C1, where the radial-priority strategy results in fewer vehicles waiting. This highlights the agent's ability to manage traffic effectively at C1 and reflects the distinct flow patterns between the circular and radial arteries.

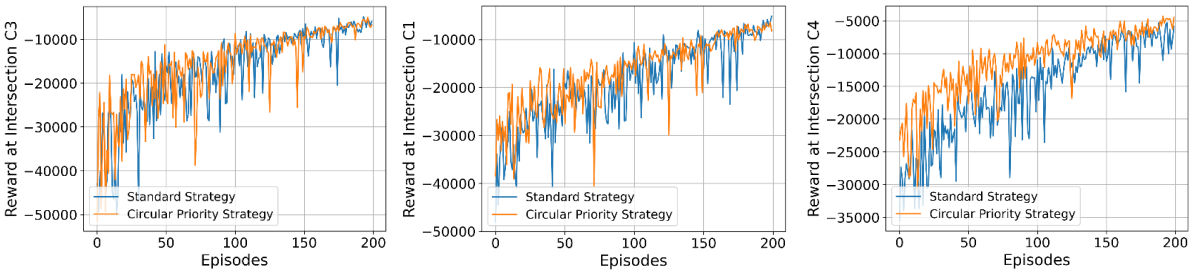

Fig. 13. Cumulative negative reward evolution for the standard and circular-priority strategies in the N>S artery.

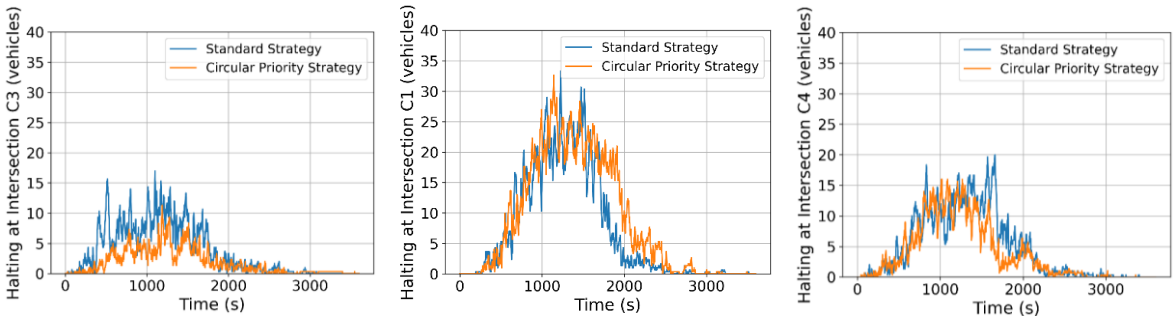

Fig. 14. Vehicle halting evolution for the standard and circular-priority strategies in the N>S artery.

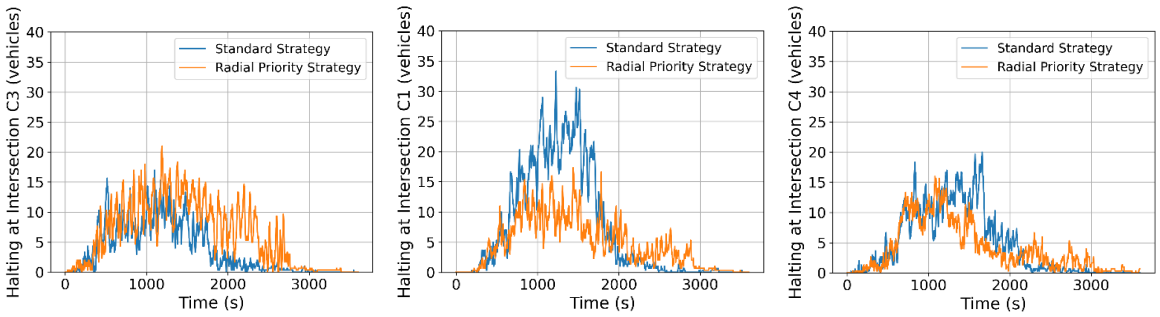

Fig. 15. Halting vehicle evolution for the standard and radial-priority strategies in the N>S artery.

Intersection C1, responsible for mediating flow between both arteries. Intersection C1, however, stands out as the most critical point in the network, as it manages traffic between both arteries. Despite the imbalance introduced in vehicle generation, the system remains stable and fluid. This can be seen in the halting values at C1, where both strategies produce very similar outcomes, demonstrating that traffic can still be effectively regulated even under asymmetric conditions.

5.4. Pedestrian Halting Analysis

Fig. 16 shows pedestrian halting across intersections for both strategies. Under the circular-priority strategy, pedestrian halting increases along the circular artery, as agents prioritize AGV movement phases. At intersections C3 and C4, however, halting is reduced due to lower vehicle density on the radial axis.

Under the radial-priority strategy, pedestrian waiting times remain generally comparable to the standard strategy but increase significantly at C1, where more frequent N–S phases reduce opportunities for pedestrian crossing. Despite this, pedestrian safety remains guaranteed, with waiting times remaining within acceptable operational limits.

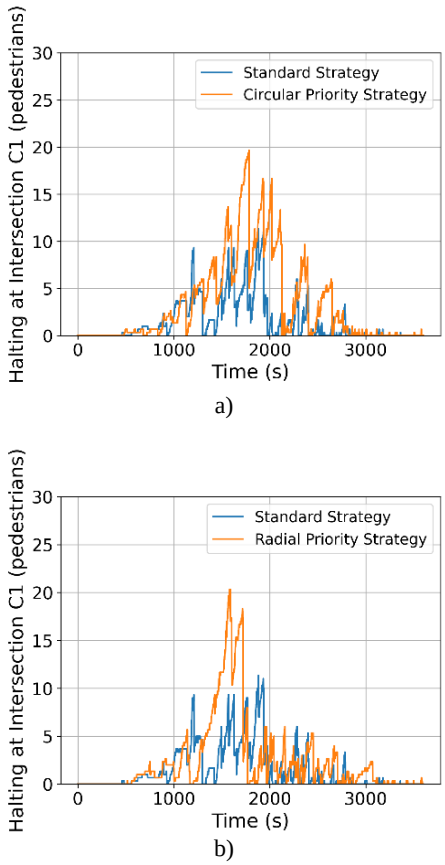

Fig. 16. Halting pedestrian evolution at C1 for a) Standard and circular-priority strategies. b) Standard and radial-priority strategies.

5.5. Activated Traffic Signal Phases

In Fig. 17 the phase-activation patterns across the standard and both circular and radial priorities are shown.

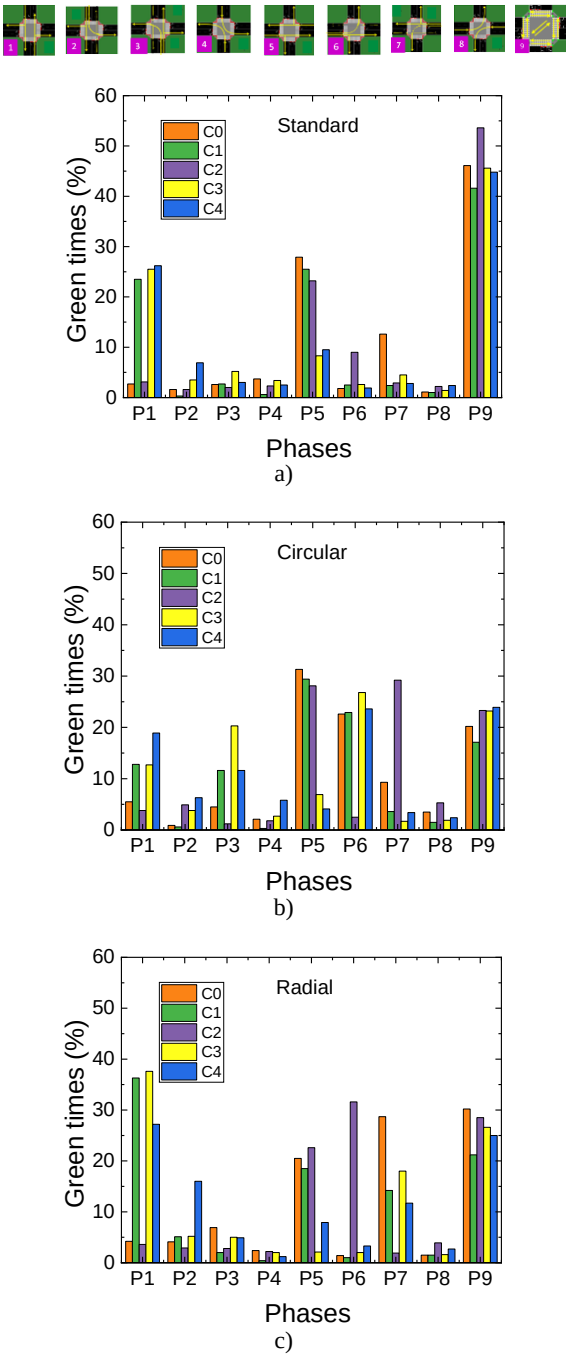

Fig. 17. Temporal comparison of green time trends across all active phases (P1-P9) at intersections C0-C4 in standard (a), circular (b) and radial (c) priority strategies.

Phase-activation patterns across strategies further confirm these trends. Under the standard strategy, balanced flows result in a corresponding equilibrium between N–S and W–E phases, with frequent pedestrian-phase activations.

The circular-priority strategy reduces N–S activations and prioritizes W–E phases – especially phases 5, 6, and 7 – at C0, C1, and C2, while significantly reducing pedestrian-phase activations. Conversely, the radial-priority strategy increases N–S phases and reduces W–E and pedestrian activations, aligning with the increased demand on the radial artery.

Each intersection also demonstrates phase preferences based on its location. C0 and C2, in the circular artery, emphasize W–E phases; C3 and C4, on the radial artery, emphasize N–S phases; and C1, linking both arteries, maintains balanced activations to stabilize overall traffic flow under varying demand conditions.

The results highlight the robustness and adaptability of the proposed DRL-based traffic management framework. Across all strategies, the agent successfully learns to adjust signal phases to accommodate local traffic demands, effectively handling both balanced and unbalanced traffic conditions. The system maintains operational stability even under strong asymmetries, particularly at intersection C1, which plays a critical role in regulating network-wide flow.

Moreover, despite variations in AGV volume, pedestrian safety and performance metrics remain within acceptable thresholds, demonstrating that multimodal traffic can be efficiently managed without compromising vulnerable users.

These findings provide a strong foundation for the broader implications discussed in the following section.

6. Conclusions

This study demonstrates the potential of integrating Visible Light Communication (VLC) with Deep Reinforcement Learning (DRL) to enable adaptive, intelligent traffic control in cyber-physical environments. VLC enhances situational awareness through precise real-time communication, while DRL dynamically adjusts signal phases to improve vehicular flow and maintain pedestrian safety.

The proposed VLC–DRL framework consistently reduces congestion, limits queue spillovers, and delivers smoother multimodal traffic behavior. Simulations show clear improvements in scalability, waiting times, and overall network fluidity, particularly in dense airport scenarios where coordination across terminals is critical.

Across all traffic strategies, the agent adapts effectively to local conditions, ensuring stable performance even under strong asymmetries. Intersection C1, the main convergence point, remains well regulated, demonstrating the system's robustness. Pedestrian safety metrics also remain within acceptable thresholds, confirming that efficiency gains do not come at the expense of vulnerable users.

Overall, the results highlight a promising path toward real-world deployment using existing lighting infrastructure. They also lay the groundwork for future extensions, including larger-scale urban networks and hybrid communication solutions to further enhance system reliability and efficiency.

Acknowledgments

This work was sponsored by FCT within the Research Unit Center of Technology and Systems CTS/UNINOVA/FCT/NOVA CTS/00066 and IPL/IDI&CA2024/INUTRAM_ISEL.

References

[1]. O. Ergul, E. Dinc, O. B. Akan, Communicate to illuminate: state-of-the-art and research challenges for visible light communications, *Physical Communication*, Vol. 17, 2015, pp. 72-85.

[2]. T. Komiyama, K. Kobayashi, K. Watanabe, T. Ohkubo, et al., Study of visible light communication system using RGB LED lights, in *Proceedings of the SICE Annual Conference*, 2011, pp. 1926-1928.

[3]. E. Monteiro, S. Hranilovic, Constellation design for color-shift keying using interior point methods, in *Proceedings of the IEEE Globecom Workshops*, 2012, pp. 1224-1228.

[4]. M. A. Vieira, M. Vieira, P. Louro, P. Vieira, et al., Vehicular visible light communication for intersection management, *Signals*, Vol. 4, Issue 2, 2023, 457.

[5]. J. Zhou, Y. Luo, H. Zhu, Deep reinforcement learning-enabled secure visible light communication against eavesdropping, *IEEE Access*, Vol. 7, 2019, pp. 103292-103301.

[6]. C. Chen, X. Shen, C. Jiang, Deep reinforcement learning-based adaptive handover mechanism for hybrid VLC and RF networks, *IEEE Transactions on Wireless Communications*, Vol. 20, Issue 11, 2021, pp. 7358-7371.

[7]. W. Shi, Y. Sun, Y. Chen, A multi-agent reinforcement learning approach with dynamic reward adjustment for traffic signal control, *IEEE Access*, Vol. 7, 2019, pp. 105186-105196.

[8]. K. Liu, Y. He, L. Li, Y. Wang, et al., Adaptive reward shaping for reinforcement learning, *IEEE Transactions on Neural Networks and Learning Systems*, Vol. 34, Issue 7, 2023, pp. 3684-3697.

[9]. M. A. Vieira, P. Louro, M. Vieira, A. Fantoni, et al., Light-activated amplification in Si-C tandem devices: a capacitive active filter model, *IEEE Sensors Journal*, Vol. 12, Issue 6, 2012, pp. 1755-1762.

[10]. M. Vieira, M. A. Vieira, P. Louro, P. Vieira, et al., Light-emitting diodes aided indoor localization using visible light communication technology, *Optical Engineering*, Vol. 57, Issue 8, 2018, 087105.

[11]. M. Vieira, M. A. Vieira, G. Galvão, P. Louro, et al., Enhancing indoor navigation in multi-terminal airports through visible light communication signals, *Proceedings of SPIE*, Vol. 12999, 2024, 129991N.

[12]. A. Alvarez Lopez, J. Blanco, F. Gonzalez, et al., Microscopic traffic simulation using SUMO, in *Proceedings of the 21st International Conference on Intelligent Transportation Systems (ITSC'18)*, 2018, pp. 2575-2582.

[13]. M. Vieira, M. A. Vieira, G. Galvão, P. Vieira, et al., AI-VLC cyber-physical system for intelligent airport traffic management, in *Proceedings of the 11th International Conference on Sensors & Electronic Instrumentation Advances (SEIA' 2025)*, 24–26 September 2025, Ponta Delgada, San Miguel, Azores Islands, Portugal, 2025, pp. 33-39.

Miguel, Azores Islands, Portugal, 2025, pp. 33-39.

[14]. M. Vieira, M. A. Vieira, P. Louro, M. Véstias, et al., Enhancing airport traffic flow: intelligent system based on VLC, rerouting techniques, and adaptive reward learning, *Sensors*, Vol. 25, 2025, 2842.