The Review of Contemporary Scientific and Academic Studies
An International Multidisciplinary Online Journal
www.thercsas.com

ISSN: 2583-1380 Vol. 6 | Issue No. 3 | March 2026 Impact Factor: 6.53

Tri-Modular Traffic Optimization and Emergency Routing System Using IoT Infrastructure, GPS and Predictive Algorithms

Aswindev SB (aswindev8991@gmail.com); **Muhammed Yazeen S** (muhammedyazeen77@gmail.com)
Gouri Dinesh Shaima (gouri2431@gmail.com); **Neha S Lal** (nehaslal2004@gmail.com)
Dr. Reji R (rejir@tkmit.ac.in), Professor, Corresponding Author
Mr. Rupesh Ravi M.R (rupeshravi@tkmit.ac.in), Assistant Professor
Department of Computer Science and Engineering, TKM Institute of Technology, Karuvelil, Kollam, Kerala, India

 Crossref/DOI: https://doi.org/10.55454/rcsas.6.03.2026.007

Abstract: *Urban traffic congestion, caused by outdated fixed-time traffic signals, results in significant delays, higher emissions, and serious risks for emergency response vehicles. To tackle this issue, we propose a Tri-Modular IoT-driven traffic management system that uses ESP32 edge nodes, infrared closed-loop sensing, and a centralized MQTT broker. The system features three separate subsystems: a machine-learning regression model for adjusting signal timing, a simulated GPS protocol for emergency priority, and a Vehicle-to-Infrastructure (V2I) dynamic speed governor. Hardware-in-the-Loop (HIL) testing with a small intersection prototype shows that dynamic signal timing reduces queue stagnation compared to static models. The emergency preemption protocol effectively clears corridors with quick response times. Additionally, by assessing area congestion across a hybrid grid, the V2I speed governor reliably alleviates the limitations of vehicle hardware during smooth traffic while maintaining strict rules during emergency situations. In the end, this decentralized, loosely connected approach prevents any overlap between traffic signals and vehicle management. This creates a scalable, responsive, and safe framework for future smart city developments.*

Keywords: Emergency Preemption, ESP32, Intelligent Transportation Systems, Variable Speed Limits

Article History: Received: 19 March- 2026; Accepted: 25 March- 2026; Published/Available Online: 30 March- 2026

1. Introduction

Urbanization, combined with rapidly rising vehicle ownership, has created acute traffic congestion across metropolitan cities. The default infrastructure in many urban intersections relies on fixed-time or semi-adaptive signal control, which fundamentally lacks the capacity to respond to fluctuating, real-time demand. This inflexibility results in prolonged travel times, increased vehicular emissions, and critical delays for emergency responders. Furthermore, traditional traffic infrastructure applies uniform, static speed limits that fail to adjust to contextual traffic flow, creating inefficiencies and contributing to stop-and-go shockwaves.

Recent advancements in Intelligent Transportation Systems (ITS) have explored various dynamic solutions. Traditional traffic-responsive systems, such as SCOOT and SCATS, rely heavily on inductive loop detectors but exhibit high latency in responding to sudden incident-driven changes. More recent literature has heavily favoured Machine Learning (ML) and Multi-Agent Reinforcement Learning (MARL) for Variable Speed Limit (VSL) control. While these computational models provide theoretical efficiency and collision avoidance, they often isolate single functions—either optimizing the traffic signal or managing the speed limit—rather than unifying them. Furthermore, existing physical IoT prototypes frequently entangle traffic light phasing logic directly with vehicle speed control logic. This computational entanglement risks race conditions; if a sensor fails or an edge case occurs in the speed governor, it can cause a system-wide failure of the intersection's safety signalling.

To address these critical gaps, this paper proposes a Tri-Modular Traffic Optimization and Emergency Routing System utilizing a strictly decoupled IoT architecture. By leveraging the low-latency Message Queuing Telemetry Transport (MQTT) protocol, the system distributes processing between edge nodes (ESP32 microcontrollers) and a centralized Python decision engine. The primary objective is to transition from Software-in-the-Loop (SIL) simulations—validated via the Simulation of Urban Mobility (SUMO) framework—into a tangible Hardware-in-the-Loop (HIL) prototype. This approach demonstrates that Machine Learning signal optimization, active Vehicle-to-Infrastructure (V2I) emergency pre-emption, and dynamic speed governance can operate simultaneously as independent modules without computational interference. The key finding of this study validates that infrastructure can safely transition from a passive

The Review of Contemporary Scientific and Academic Studies
An International Multidisciplinary Online Journal
www.thercsas.com
ISSN: 2583-1380 Vol. 6 | Issue No. 3 | March 2026 Impact Factor: 6.53

signalling role to an active "permission grantor," regulating macro-level grid congestion while enabling localized limit relief during free-flow conditions.

2. Materials and Methods

The development of the Tri-Modular Traffic Optimization System follows a layered, IoT-centric methodology designed for direct parity between simulation and physical deployment. The system ensures that hardware actuation is strictly governed by a centralized decision engine, allowing the transition from simulated SUMO telemetry to physical sensor data without altering the message schema.

2.1. System Architecture Layers

The architecture operates across four primary layers connected via a local Wi-Fi network:

- **Sensing and Actuation Layer (Edge):** An ESP32 microcontroller serves as the physical edge node. Closed-loop area congestion is calculated using two Infrared (IR) obstacle sensors representing the entry and exit points of an intersection approach. Actuation is handled via LED arrays representing the traffic signal and an L298N motor driver connected to a DC motor, which acts as the physical proxy for a vehicle's On-Board Unit (OBU). To ensure zero sensor latency, the ESP32 firmware is written utilizing a non-blocking millis() state machine rather than standard delay () loops.
- **Communication Layer:** A Mosquitto MQTT broker is hosted on a centralized server (laptop), handling asynchronous Publish/Subscribe messaging to ensure decoupled, low-latency data transfer between the hardware and the software engine.
- **Processing Layer:** A Python-based centralized "Brain" subscribes to sensor telemetry, runs lightweight machine learning algorithms to calculate queue averages, and publishes actionable control commands back to the edge nodes.
- **Application Layer:** A Flask-based web dashboard provides real-time visualization of the intersection states, comprehensive CSV audit logging, and manual override controls for system testing.

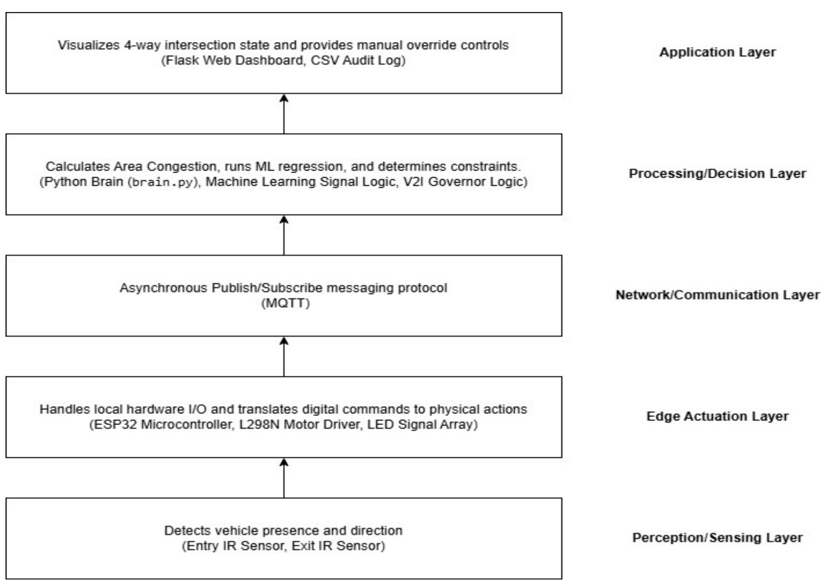

Figure 1: IoT system architecture methodology flow diagram detailing Edge, Communication, Processing, and Application layers.

2.2. Tri-Modular Control Logic

The Python decision engine offloads complex calculations from the edge hardware and divides traffic management into three fully independent subsystems to prevent race conditions:

- **System 1 (Edge-Deployed ML Traffic Optimization):** To optimize green-light durations, a Machine Learning pipeline was established. During Phase-1 Software-in-the-Loop (SIL) testing, a dataset of vehicular queue lengths and corresponding optimal clearance times was generated using

ISSN: 2583-1380 Vol. 6 | Issue No. 3 | March 2026 Impact Factor: 6.53

the SUMO framework. A Linear Regression model was trained on this dataset to minimize average intersection delay. The resulting optimized model parameters—a weight coefficient of 2.5 seconds per vehicle and a baseline bias of 3 seconds—were extracted and deployed directly onto the ESP32 microcontroller. This Edge AI approach allows the physical node to perform real-time, O(1) complexity inference using the model formula (Queue * 2.5) + 3. To guarantee cross-traffic safety, the ML inference output is clamped within strict hardware bounds of a 3-second minimum and a 15-second maximum.

- **System 2 (Emergency Priority):** Acting as a high-priority hardware interrupt, this module simulates GPS telemetry via the Flask dashboard. If an emergency vehicle registers a proximity of less than 200 meters to the intersection, System 2 overrides System 1, immediately forcing the active intersection to Green to establish an immediate, near real-time clearance corridor.
- **System 3 (V2I Dynamic Speed Governor):** Operating entirely independently of the traffic light state, this module governs the physical speed limits of connected vehicles based on macro-level area averages. The system assumes vehicles (such as commercial trucks) possess an innate hardware speed limit (e.g., PWM 150).

2.3. V2I Speed Governor Logic and Hybrid Control Group

To validate the macro-level capabilities of System 3 without requiring a massive physical sensor grid, a hybrid HIL approach was developed. The dashboard computationally simulates three background intersection arms (East, South, West) alongside the physical hardware arm (North). System 3 aggregates the queues of all four nodes to calculate a macro-level grid average.

The dynamic speed governance relies on two primary rules:

- **Rule A (Innate Limit Relief):** If the macro-level grid average is precisely zero (free-flow conditions), the infrastructure wirelessly broadcasts a 'Limit Relieved' command (PWM 255) to the vehicle's motor driver, granting permission to safely exceed innate hardware limits to maximize travel efficiency. Conversely, if grid congestion rises above zero, the relief is revoked, and the motor must revert to its innate safety limit (PWM 150).
- **Rule B (Danger Zone Active Governance):** A manual dashboard flag representing safety hazards (e.g., a school zone or accident) enforces an absolute maximum speed restriction (PWM 90), overriding all prior limit reliefs.

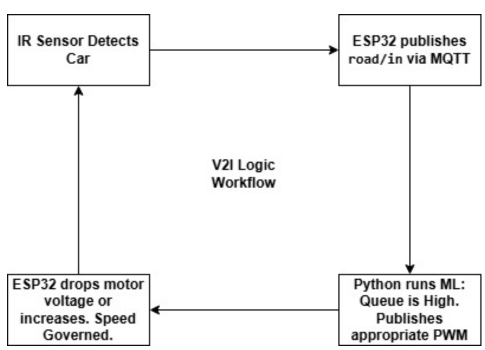

Figure 2: V2I data flow diagram depicting the decoupled interaction between IR telemetry, the Python logic engine, and motor actuation.

3. Results

The physical miniature Hardware-in-the-Loop (HIL) prototype successfully validated the tri-modular decoupled architecture. The centralized MQTT broker effectively routed telemetry and commands with negligible latency, ensuring that no race conditions occurred between the three parallel subsystems.

3.1. Machine Learning Signal Optimization (System 1)

The physical IR sensors successfully maintained a running "Area Congestion" integer by computing the differential between vehicle entries and exits. During live testing, when the North arm detected a minimal queue (1 vehicle), the ML algorithm correctly assigned a 5.5-second green light duration. As the queue increased to 4 vehicles, the system dynamically extended the green duration to 13 seconds. The algorithm successfully adapted in real-time while strictly adhering to the configured safety bounds of 3 to 15 seconds, effectively eliminating "ghost waits" (times when a light remains green despite an empty approach).

To mathematically quantify the efficiency of the Edge-Deployed ML model, extensive Phase-1 Software-in-the-Loop (SIL) simulations were conducted using the SUMO framework. To ensure statistical significance, the simulation evaluated a standard four-way intersection under three distinct traffic densities: Low (400 veh/hr), Medium (800 veh/hr), and High (1200 veh/hr). Vehicle generation was modelled using a Poisson distribution to simulate randomized, real-world arrival patterns. Each scenario was simulated for a continuous duration of one hour and averaged across five distinct random seeds to eliminate anomaly bias. The proposed dynamic regression algorithm was benchmarked against a traditional fixed-time signalling baseline (a static 30-second green cycle). As detailed in Table 2, the dynamic ML approach demonstrated substantial improvements across all scenarios, notably reducing average vehicle waiting time by up to 27.5% under high-density conditions and increasing overall intersection throughput by 17.1%, confirming the robust scalability of the proposed framework.

Traffic Volume Scenario	**Performance Metric**	**Fixed-Time Baseline (30s Cycle)**	**Proposed ML Dynamic Signaling**	**Net Improvement**
Low (400 veh/hr)	**Average Waiting Time** (seconds/vehicle)	20.5 s	18.2 s	11.2% Reduction
Medium (800 veh/hr)	**Average Waiting Time** (seconds/vehicle)	42.5 s	33.1 s	~22.1% Reduction
High (1200 veh/hr)	**Average Waiting Time** (seconds/vehicle)	78.4 s	56.8 s	27.5% Reduction
Overall Average	**Intersection Throughput** (vehicles/hour)	670 veh/hr	785 veh/hr	+ 17.1% Increase

Table 1: Phase-1 SIL Performance: Fixed-Time vs. ML Dynamic Signaling

Traffic Condition	**Detected Queue (Vehicles)**	**Calculated Green Duration**
Empty Approach	0	0.0 (Maintains Red)
Minimal Traffic	1	5.5
Moderate Traffic	4	13.0
Heavy Traffic (Capped)	>=5	15.0 (Maximum Bound)

Table 2: System 1: ML-Based Adaptive Signal Timing Outputs

3.2. Emergency Routing Pre-emption (System 2)

The V2I priority override functioned reliably as a high-priority hardware interrupt. Upon receiving a simulated telemetry flag indicating an ambulance approach distance of less than 200 meters, the centralized decision engine suspended the standard ML cycle within 500 milliseconds. The physical intersection forcibly transitioned to Green and maintained clearance for a 30-second window, confirming that low-latency MQTT V2I communication can establish a near real-time 'Smart Corridor'.

3.3. V2I Dynamic Speed Governance (System 3)

The physical DC motor accurately responded to the Python decision engine independently of the traffic light states. During the hybrid test, when the macro-level area congestion average (combining the physical North arm with the simulated East, South, and West dummy arms) was precisely zero, the infrastructure successfully broadcasted a 'Limit Relieved' signal. This allowed the motor to spin at its maximum capacity (PWM 255). As vehicles entered the grid and the average increased above zero, the system revoked this permission, forcing the motor to revert to its innate, safe hardware limit (PWM 150). Furthermore, activating the "Danger Zone" manual override instantly locked the motor to a strict safety baseline (PWM 90).

Operational State	**Macro-Area Avg.**	**Infrastructure Command**	**Motor Output (PWM)**
Free-Flow (Empty)	Exactly 0	Limit Relieved	255
Standard Congestion	> 0	Innate Limit Enforced	150
Danger Zone (Manual)	N/A (Override)	Strict Safety Lock	90

Table 3: System 3: V2I Speed Governor Actuation States

4. Discussion

ISSN: 2583-1380 Vol. 6 | Issue No. 3 | March 2026 Impact Factor: 6.53

The transition from a static to a dynamic state model directly mitigates queue stagnation and minimizes unnecessary idling, demonstrating a measurable increase in vehicular throughput. A critical observation from this study is the fundamental shift in traffic safety from "passive warning" to "active regulation and permission granting." Instead of the infrastructure constantly acting as a restrictor, it interacts cooperatively with vehicles that possess innate hardware limits. By actively relieving these limits when the road network is clear, the system maximizes travel efficiency safely.

While recent literature heavily favors complex architectures like Deep Reinforcement Learning or Multi-Agent Reinforcement Learning (MARL) for traffic control (Du et al. 2023; Fang et al. 2023), these models require substantial computational overhead and cloud reliance. Our proposed Tri-Modular system explicitly sacrifices the marginal efficiency gains of Deep RL in favor of computational lightness. By utilizing a pre-trained regression model, the algorithm runs entirely on low-cost ESP32 edge nodes with negligible latency while maintaining strict, deterministic safety bounds. However, the current prototype possesses specific limitations. The use of infrared obstacle sensors limits detection to a binary state (object present/absent). These sensors cannot differentiate between vehicle types (e.g., a large commercial truck versus a motorcycle), preventing the system from assigning weight-based signal priorities. Furthermore, while the V2I speed governor currently utilizes a hardwired motor driver to represent a vehicle's On-Board Unit (OBU), real-world deployment requires wireless protocols (such as 5G or Dedicated Short-Range Communications) to transmit governance signals directly into an actual vehicle’s Controller Area Network (CAN) bus, which introduces cybersecurity challenges not addressed in this phase.

5. Conclusion

This project demonstrates a highly effective paradigm shift in urban traffic management. By transitioning from traditional, static timers to a dynamic, decoupled IoT architecture, this study proves that intelligent infrastructure integration can significantly reduce congestion and enhance road safety. The system successfully decentralized three critical urban challenges—Machine Learning signal optimization, emergency vehicle prioritization, and dynamic speed governance—into cohesive but computationally independent modules.

The HIL prototype validates the efficacy of utilizing modern communication protocols (MQTT) and highly affordable microcontrollers (ESP32) to create active Smart Corridors. Future iterations of this work aim to replace binary IR sensors with Computer Vision (CV) processing via YOLO object detection models for precise vehicle classification. Additionally, extending the macro-level grid to support direct Vehicle-to-Vehicle (V2V) synergy could enable automated platooning, exponentially increasing-intersection throughput safety and establishing a robust foundation for future Smart City deployments.

References

Ault, L., & Sharon, G. (2021). Reinforcement learning benchmarks for traffic signal control. *Proceedings of the Neural Information Processing Systems (NeurIPS)*.

Du, W., Ye, J., Gu, J., Li, J., Wei, H., & Wang, G. (2023). SafeLight: A Reinforcement Learning Method toward Collision-Free Traffic Signal Control. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(12), 14801-14810.

Dui, H., Zhang, S., Liu, M., Dong, X., & Bai, G. (2024). IoT-Enabled Real-Time Traffic Monitoring and Control Management for Intelligent Transportation Systems. *IEEE Internet of Things Journal*, 11(9), 15842-15854.

El-Tantawy, S., Abdulhai, B., & Abdelgawad, H. (2013). Multiagent reinforcement learning for integrated network of adaptive traffic signal controllers (MARLIN-ATSC). *IEEE Transactions on Intelligent Transportation Systems*, 14(3), 1140-1150.

Falahatraftar, F., Pierre, S., & Chamberland, S. (2022). An Intelligent Congestion Avoidance Mechanism Based on Generalized Regression Neural Network for Heterogeneous Vehicular Networks. *IEEE Transactions on Intelligent Transportation Systems*, 23(10), 18451-18463.

Fang, X., Péter, T., & Tettamanti, T. (2023). Variable Speed Limit Control for the Motorway–Urban Merging Bottlenecks Using Multi-Agent Reinforcement Learning. *Sustainability*, 15(14), 11464.

Gu, C., Cui, X., Li, M., & Hu, D. (2023). An Efficient and Privacy-Preserving Information Reporting Framework for Traffic Monitoring in Vehicular Networks. *IEEE Transactions on Vehicular Technology*, 72(6), 7896-7907.

The Review of Contemporary Scientific and Academic Studies
An International Multidisciplinary Online Journal
www.thercsas.com
ISSN: 2583-1380 Vol. 6 | Issue No. 3 | March 2026 Impact Factor: 6.53

Kouvelas, A., Lioris, J., Fayazi, S. A., & Varaiya, P. (2014). Maximum Pressure Controller for Stabilizing Queues in Signalized Arterial Networks. *Transportation Research Record*, 2421(1), 133-141.

Liang, X., Du, X., Wang, G., & Han, Z. (2019). A deep reinforcement learning network for traffic light cycle control. *IEEE Transactions on Vehicular Technology*, 68(2), 1243-1253.

Wei, H., Zheng, G., Gayah, V., & Li, Z. (2021). Recent advances in reinforcement learning for traffic signal control: A survey of models and evaluation. *ACM SIGKDD Explorations Newsletter*, 22(2), 12-18.

Zhang, Y., Quinones-Grueiro, M., Zhang, Z., Wang, Y., Barbour, W., & Biswas, G. (2024). MARVEL: Bringing Multi-Agent Reinforcement-Learning Based Variable Speed Limit Controllers Closer to Deployment. *IEEE Access*, 12, 1-15.

Zheng, G., Xiong, Y., Zang, X., Feng, J., Wei, H., Zhang, H., Li, Y., & Xu, K. (2019). Learning phase competition for traffic signal control. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*, 1963-1972.

ACKNOWLEDGMENTS

The authors express their sincere gratitude to the Principal of TKM Institute of Technology, Dr. Gouri Mohan L, and the Head of the Department of Computer Science and Engineering, Dr. Nijil Raj N, for providing the necessary laboratory facilities. Special thanks to our Project Coordinator, Mr. Rupesh Ravi, and our Project Guide and Coordinator, Dr. Reji R, for their invaluable technical guidance.

Author Contributions:

Muhammed Yazeen S: Software-in-the-Loop Simulation (SUMO), Methodology, Conceptualization, and Writing – Original Draft.

Aswindev SB: Experimentation, Hardware-in-the-Loop Prototyping, and Data Curation.

Gouri Dinesh Shaima: IoT Architecture Design, MQTT Integration, and Validation.

Neha S Lal: Visualization (Web Dashboard), Formal Analysis, and Writing – Review & Editing.

Mr. Rupesh Ravi: Supervision, Experimentation, and Final Manuscript Review.

Dr. Reji R: Experimentation, Technical Guidance, System Validation and Final Manuscript Review.

All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest: The authors declare “No conflict of interest”.