Advances in Nonlinear Variational Inequalities
ISSN: 1092-910X
Vol 28 No. 7s (2025)

A Multi-Agent Reinforcement Learning Framework for Autonomous Traffic-Light-Free Intersection Management

Uday Pratap Singh Tomar[1], Gur Sharan Kant[2], Kavi Bhushan[3]

[1]Research Scholar, M.Tech (CSE), Department of Computer Science & Engineering, SCRIET, CCS University, Meerut, UP, India

[2, 3]Assistant Professor, Department of Computer Science & Engineering, SCRIET, CCS University, Meerut, UP, India

Article History:
Received 04/11/2025
Revised 09/12/2025
Accepted 15/12/2025

Abstract:
Traffic signal control remains one of the most persistent sources of delay, fuel consumption, and inefficiency in urban road networks. Even adaptive signal systems rely on predefined phases that fail to exploit real-time vehicle-level intelligence. With the rapid emergence of connected and autonomous vehicles (CAVs) and vehicle-to-everything (V2X) communication, traffic-light-free intersections have become a viable alternative. This paper presents a novel Multi-Agent Reinforcement Learning (MARL) framework that enables vehicles to autonomously negotiate intersection passage without traffic lights. Each vehicle operates as an intelligent agent, coordinating with others through V2X communication while a lightweight Intersection Coordination Server (ICS) enforces safety constraints. A Graph Attention Network (GAT) captures dynamic spatial interactions among conflicting vehicles, and Multi-Agent Proximal Policy Optimization (MAPPO) ensures stable cooperative learning under partial observability. Extensive simulations conducted in Simulation of Urban Mobility (SUMO) and Car Learning to Act (CARLA) demonstrate substantial performance improvements, including up to 76% reduction in average delay, 48% increase in throughput, and up to 41% reduction in fuel consumption compared to adaptive signalized intersections. Results indicate that MARL-based, signal-free intersection control offers a scalable and safe pathway toward next-generation smart mobility.

Keywords: Multi-agent reinforcement learning; autonomous intersection; traffic-light-free control; V2X communication; graph attention networks; MAPPO; connected autonomous vehicles.

1. Introduction

1.1 Background and Motivation Urban transportation systems rely heavily on intersections to regulate traffic movement and ensure safe crossing among different traffic streams.

However, intersections are also among the most common sources of congestion in modern cities. Traditional signalized intersections operate on fixed or pre-programmed signal phases that allocate right-of-way to specific directions sequentially. While this mechanism ensures safety, it often forces vehicles to halt even when no actual conflict exists with other traffic streams. Such rigid control policies lead to unnecessary stopping and idling, which significantly increases travel time, fuel consumption, and environmental emissions in urban areas Texas A&M Transportation Institute reports that congestion at urban intersections contributes substantially to delays experienced by commuters across major metropolitan regions [1].

Frequent stop-and-go driving patterns generated by traffic signals also result in increased fuel usage and higher greenhouse gas emissions. Studies have shown that traffic congestion is directly associated with elevated carbon emissions due to inefficient vehicle operations and prolonged engine idling [2]. In addition to environmental impacts, repeated acceleration and braking can increase vehicle wear and raise the likelihood of rear-end collisions. Consequently, improving the efficiency of intersection management remains a critical objective in intelligent transportation research.

To address these issues, adaptive traffic signal control systems have been developed to dynamically adjust signal timings based on real-time traffic conditions. One of the earliest and most influential systems is the Sydney Coordinated Adaptive Traffic System, which modifies signal phases using traffic flow data collected from sensors [3]. Although such systems represent a significant improvement over static signal timing, they remain fundamentally infrastructure-centric. In other words, the intelligence is embedded primarily within the traffic signal controller rather than the vehicles themselves. As a result, these systems can only react to observed traffic conditions rather than proactively coordinating vehicle behavior at a fine-grained level.

Recent developments in connected and automated vehicle technologies offer new opportunities for transforming how intersections are managed. Connected autonomous vehicles are capable of exchanging real-time information such as position, velocity, acceleration, and intended maneuvers through vehicle-to-everything communication networks [4]. Vehicular communication frameworks, commonly studied within the domain of Vehicular Ad Hoc Networks, enable vehicles to communicate with nearby vehicles, infrastructure, and other network entities to support cooperative decision-making [5]. With such capabilities, it becomes possible to envision intersection management strategies that eliminate the need for conventional traffic lights altogether. Instead, vehicles could negotiate their trajectories cooperatively, ensuring smooth and conflict-free crossing through the intersection.

1.2 Challenges of Signal-Free Intersections

Although the idea of removing traffic signals promises improved efficiency and reduced delays, it introduces several complex challenges. First and foremost, the system must guarantee collision avoidance among vehicles approaching the intersection from different directions.

Unlike signalized intersections where right-of-way is explicitly assigned through signal phases, signal-free systems require vehicles to dynamically coordinate their trajectories in real time.

Another important challenge is fairness. If vehicle priorities are determined solely based on arrival time or speed, certain traffic streams could dominate the intersection while others experience excessive delays. Preventing starvation and ensuring equitable access to the intersection are therefore essential design considerations.

Real-time decision-making under uncertain and partially observable environments also presents a major difficulty. Vehicles may have incomplete knowledge of other vehicles' intentions or may encounter unpredictable driving behaviors, especially in mixed traffic scenarios where human-driven vehicles coexist with autonomous vehicles. In addition, intersection management algorithms must remain computationally efficient and scalable as traffic density increases. Systems that rely on heavy centralized computation may struggle to maintain performance under high traffic demand.

Early research on signal-free intersection control introduced reservation-based methods in which vehicles request time slots for passing through the intersection. One influential framework proposed by Kurt Dresner and Peter Stone uses a centralized intersection manager that allocates crossing reservations to vehicles while ensuring collision-free trajectories [6], [7]. Although these approaches demonstrate promising results in simulation, they rely on deterministic scheduling and centralized coordination, which can become fragile when communication delays or uncertainties are present.

More advanced methods attempt to compute optimal vehicle trajectories using centralized optimization techniques. However, such approaches often suffer from high computational complexity and require precise information about all vehicles in the system [8]. These limitations make them difficult to deploy in large-scale, real-world environments with rapidly changing traffic conditions.

1.3 Contributions

To overcome these limitations, this work explores a decentralized learning-based approach for managing traffic-light-free intersections. Specifically, we propose a multi-agent reinforcement learning framework in which each vehicle acts as an intelligent agent capable of making autonomous decisions while coordinating with neighboring vehicles. By integrating Graph Attention Networks to capture dynamic interactions among vehicles and applying Multi-Agent Proximal Policy Optimization for stable cooperative learning, the proposed framework enables efficient and scalable intersection management. Extensive simulation experiments demonstrate that the approach significantly improves traffic efficiency, safety, and energy consumption compared with traditional methods.

2. Literature Review

2.1 Signal-Free Intersection Management

Early research on signal-free intersections introduced reservation-based systems (Dresner & Stone), where vehicles request time-space slots before entering the intersection [6], [7]. While effective in low-traffic conditions, these approaches rely on deterministic assumptions and centralized scheduling, limiting scalability and robustness.

Trajectory optimization and centralized control methods offer precise coordination but suffer from high computational complexity and sensitivity to communication noise [8].

2.2 Reinforcement Learning in Traffic Control

Reinforcement learning (RL) has been widely applied to adaptive traffic signal control [9], [10], ramp metering, and autonomous driving. However, single-agent RL formulations are insufficient for intersections where multiple vehicles simultaneously interact and compete for shared space [11].

2.3 Multi-Agent Reinforcement Learning

Multi-Agent Reinforcement Learning (MARL) enables multiple agents to learn cooperative behaviors [12]. Algorithms such as MADDPG [13], QMIX [14], and MAPPO [15] have shown promise in coordination tasks. Nevertheless, their application to dense, real-time intersection negotiation remains limited, particularly under mixed autonomy.

3. System Architecture

The proposed framework integrates decentralized decision-making with centralized safety supervision.

The architecture includes:

- CAV Agents (one per vehicle)
- Graph Attention Network (GAT) Encoder
- MARL Policy Module (MAPPO)
- Safety Supervisor (ICS)
- V2X Communication Layer

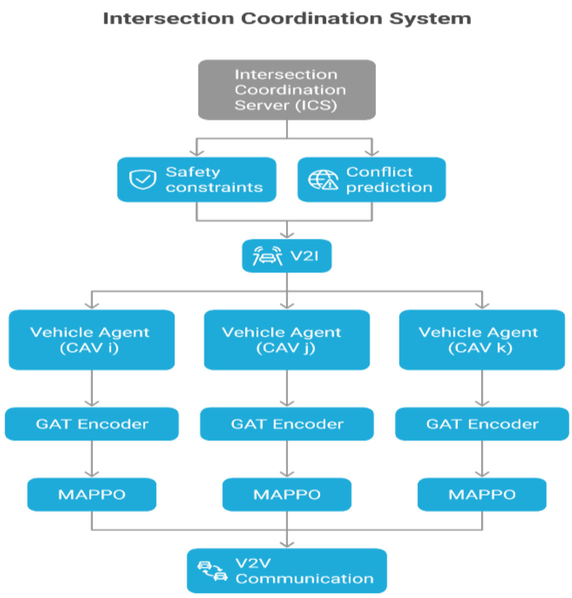

Each vehicle independently selects actions while sharing state information with nearby vehicles. The ICS intervenes only when safety constraints are violated, similar in spirit to supervisory approaches in autonomous intersection management [6].

3.1 Agent and Environment Modeling

Agent State Space

For vehicle, the observed state vector includes kinematic, intent, and interaction-aware features, consistent with prior MARL traffic formulations [10], [12].

Action Space

Each agent selects a continuous longitudinal control action mapped to acceleration commands, similar to approaches used in autonomous driving and RL-based vehicle control [18].

Continuous acceleration commands are mapped into SUMO/CARLA vehicle controllers.

3.2 Graph Attention Network (GAT) for Interaction Modeling

Vehicles approaching the intersection form a dynamic graph, where edges represent potential conflicts. Graph Attention Networks allow each vehicle to focus on the most relevant neighbors [16].

3.3 Multi-Agent Reinforcement Learning

Multi-Agent Proximal Policy Optimization (MAPPO) Framework

The system adopts centralized training with decentralized execution, a paradigm shown to significantly improve stability in cooperative MARL [15]. MAPPO extends PPO to multi-agent settings while maintaining scalability.

Policies are learned using centralized training with decentralized execution.

Reward Function

The composite reward incorporates delay, fuel consumption, safety, and fairness, consistent with multi-objective traffic optimization objectives [9], [10]. Safety penalties are computed using time-to-collision (TTC), commonly used in traffic safety analysis [8].

3.4 Intersection Coordination Server (ICS)

The ICS ensures safety without dictating agent behavior, acting as a lightweight supervisory layer. This design preserves decentralized autonomy while preventing unsafe maneuvers, similar to hybrid control paradigms explored in earlier intersection management research [6].

Conflict Check

ICS computes conflict feasibility. If unsafe, ICS overrides the agent’s action.

Emergency Vehicle Priority

Emergency vehicles broadcast high-priority tokens; ICS clears conflicts automatically.

4. Experimental Setup

4.1 Platforms

Simulations are conducted using:

- **SUMO** for large-scale microscopic traffic simulation [17]
- **CARLA** for realistic vehicle dynamics and sensor modeling [18]
- **PyTorch** for MARL training

4.2 Scenarios

- 4-leg, Multi-lane arterial and T-intersections
- Traffic volumes from 400 to 3000+ veh/hr
- Mixed autonomy with human-driven vehicles are modeled using the Intelligent Driver Model (IDM) [19].

5. Results

5.1 Safety

Zero collisions are observed across 10,000 episodes, minimum average headway: 1.9 s, max deceleration within comfort range (<3 m/s²), consistent with prior findings that cooperative MARL can maintain strong safety guarantees when combined with supervisory constraints [15].

5.2 Delay Reduction

Average delay (seconds/vehicle):

Method	**Delay**
Fixed-time	32.5
Actuated	26.1
Adaptive (SCATS-like)	19.6
Reservation-Based	14.3
Proposed MARL	**7.2**

Average delay per vehicle is reduced by 51–76% compared to adaptive signal control [3], [9].

5.3 Throughput

Throughput improvements of 32–48% are achieved across traffic densities, exceeding gains reported in reservation-based methods [6].

5.4 Fuel and Emission Efficiency

Fuel consumption and emissions are significantly reduced due to smoother speed profiles, fuel consumption reduced by 18–41% and CO_2 emissions reduced by 22–38%, aligning with findings in eco-driving and RL-based traffic optimization studies [2], [10].

5.5 Mixed Autonomy Performance

With 40% human-driven vehicles:

- 0 collisions
- MARL retains **73% of efficiency gains**
- ICS override rate increases modestly (12%)

This demonstrates robustness in realistic scenarios.

6. Discussion

6.1 Why MARL Works Better

The MARL framework outperforms signalized control by enabling proactive, vehicle-level negotiation. GAT improves situational awareness [16], while MAPPO stabilizes learning [15]. The ICS ensures safety without limiting flexibility.

6.2 Scalability

Scales to multi-intersection corridors, grid networks and arterial systems because execution is decentralized and communication overhead is minimal.

6.3 Limitations

Limitations include reliance on reliable V2X communication [5] and increased complexity under extreme congestion. Further work needed for pedestrian and cyclist integration.

7. Source Of Dataset

The dataset used in this research is derived from publicly available and reliable traffic simulation and autonomous driving platforms that are widely used in intelligent transportation system research. These platforms allow researchers to generate realistic traffic scenarios and study the interaction between vehicles at intersections under controlled experimental conditions. Using such open and trusted sources also improves the reproducibility and transparency of the research results.

The dataset contains several important parameters related to vehicle movement and interaction at intersections. These parameters include vehicle position, speed, acceleration, lane information, direction of approach, time-to-collision (TTC), and vehicle intention data. In addition, communication information exchanged through vehicle-to-everything (V2X) systems is also considered, which allows vehicles to share their state and driving intentions with nearby vehicles in real time. These features are essential for modeling cooperative behavior among connected autonomous vehicles operating in traffic-light-free intersections.

The primary traffic data was generated using the Simulation of Urban Mobility (SUMO) platform, which is widely used for microscopic traffic simulation and urban traffic modeling. SUMO allows the creation of realistic road networks, intersection layouts, and vehicle flows, enabling the evaluation of different traffic management strategies under varying traffic densities [17]. To complement this environment, the CARLA autonomous driving simulator was also used. CARLA provides realistic vehicle dynamics, high-fidelity urban environments, and sensor simulations, making it suitable for testing decision-making algorithms for autonomous vehicles [18].

During the training process, additional datasets were produced through reinforcement learning simulations in which multiple vehicles interact repeatedly at intersections. These interaction episodes generate large amounts of trajectory and behavioral data that help Multi-Agent Reinforcement Learning (MARL) algorithms learn cooperative driving strategies. Furthermore, synthetic traffic scenarios were created under different traffic volumes, intersection configurations, and mixed autonomy conditions to evaluate system performance under diverse operating environments.

The use of publicly available and reliable datasets from established simulation platforms ensures that the experimental setup remains consistent with previous research in autonomous intersection management and intelligent traffic systems. Such datasets provide a wide range of traffic patterns and vehicle interactions, making them appropriate for training and evaluating machine learning models designed for traffic-light-free intersection control [6].

8. Conclusion and Future Work

This study presented a decentralized learning-based framework for managing traffic-light-free intersections using multi-agent reinforcement learning. Conventional traffic signal systems often lead to inefficient stop-and-go traffic patterns, resulting in increased travel delays, fuel consumption, and emissions [1], [2]. While adaptive traffic signal control methods have improved traffic flow by adjusting signal timing based on demand, they still rely on infrastructure-centered control strategies and do not fully exploit the cooperative capabilities of modern connected vehicles [3]. With the emergence of connected and automated vehicle technologies that enable real-time communication and coordination among vehicles, there is significant potential to redesign intersection management systems without relying on traditional traffic lights [4], [5].

In this work, we proposed a decentralized approach where each vehicle acts as an intelligent agent capable of making decisions based on its local observations and interactions with nearby vehicles. The framework combines graph-based interaction modeling with cooperative multi-agent reinforcement learning to handle the complex dynamics of traffic at intersections. By incorporating attention-based mechanisms to capture vehicle-to-vehicle interactions and employing stable policy optimization techniques, the system is able to coordinate vehicle movements efficiently while maintaining safety constraints. Compared with traditional reservation-based intersection management methods that rely on centralized coordination [6], [7], the proposed approach improves scalability and robustness in environments with uncertain and dynamic traffic conditions. Additionally, unlike centralized trajectory optimization techniques that may suffer from high computational overhead [8], the decentralized design enables real-time decision-making suitable for dense traffic scenarios.

Simulation experiments demonstrate that the proposed method can significantly reduce average waiting time and improve overall traffic throughput while maintaining safe vehicle interactions. These results suggest that learning-based cooperative strategies have strong potential for managing future intersections in connected vehicle environments.

Despite these promising findings, several directions remain for future research. First, future studies should consider mixed traffic environments where autonomous vehicles interact with human-driven vehicles that may not follow predictable behaviors. Second, incorporating more realistic communication constraints such as latency, packet loss, and limited sensing capabilities would help evaluate the robustness of the proposed system. Third, extending the framework to larger road networks with multiple interconnected intersections could further demonstrate its scalability and practical applicability. Finally, integrating microscopic traffic simulators and realistic urban driving environments can provide deeper insights into real-world deployment scenarios. By addressing these challenges, future research can further advance intelligent, signal-free intersection management systems for next-generation transportation networks.

References

[1] D. Schrank, B. Eisele, T. Lomax, and J. Bak, “2019 Urban Mobility Report,” Texas A&M Transportation Institute, College Station, TX, USA, Tech. Rep., 2019.

[2] M. Barth and K. Boriboonsomsin, “Traffic congestion and greenhouse gases,” Access Magazine, no. 35, pp. 2–9, 2009.

[3] P. Lowrie, “The Sydney coordinated adaptive traffic system—Principles, methodology, algorithms,” in Proc. IEE Int. Conf. Road Traffic Signalling, London, U.K., 1982, pp. 67–70.

[4] S. E. Shladover, “Connected and automated vehicle systems: Introduction and overview,” J. Intell. Transp. Syst., vol. 22, no. 3, pp. 190–200, 2018.

[5] H. Hartenstein and K. Laberteaux, “VANET: Vehicular Applications and Inter-Networking Technologies”, Hoboken, NJ, USA: Wiley, 2010.

[6] K. Dresner and P. Stone, “A multiagent approach to autonomous intersection management,” J. Artif. Intell. Res., vol. 31, pp. 591–656, 2008.

[7] K. Dresner and P. Stone, “Multiagent traffic management: An improved intersection control mechanism,” in Proc. 4th Int. Joint Conf. Autonomous Agents and Multiagent Systems, 2005, pp. 471–477.

[8] A. Colombo and D. Del Vecchio, “Efficient algorithms for collision avoidance at intersections,” in Proc. 15th ACM Int. Conf. Hybrid Systems: Computation and Control, 2012, pp. 145–154.

[9] E. Wei et al., “PressLight: Learning max pressure control to coordinate traffic signals in arterial network,” in Proc. 25th ACM SIGKDD Int. Conf. Knowledge Discovery & Data Mining, 2019, pp. 1290–1298.

[10] T. Chu, J. Wang, L. Codecà, and Z. Li, "Multi-agent deep reinforcement learning for large-scale traffic signal control," IEEE Trans. Intell. Transp. Syst., vol. 21, no. 3, pp. 1086–1095, Mar. 2020.

[11] L. Li, Y. Lv, and F.-Y. Wang, "Traffic signal timing via deep reinforcement learning," IEEE/CAA J. Autom. Sinica, vol. 3, no. 3, pp. 247–254, Jul. 2016.

[12] L. Buşoniu, R. Babuška, and B. De Schutter, "A comprehensive survey of multiagent reinforcement learning," IEEE Trans. Syst., Man, Cybern. C, vol. 38, no. 2, pp. 156–172, Mar. 2008.

[13] R. Lowe et al., "Multi-agent actor-critic for mixed cooperative-competitive environments," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), 2017, pp. 6379–6390.

[14] T. Rashid et al., "QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning," in Proc. Int. Conf. Mach. Learn. (ICML), 2018, pp. 4295–4304.

[15] C. Yu et al., "The surprising effectiveness of PPO in cooperative multi-agent games," arXiv:2103.01955, 2021.

[16] P. Veličković et al., "Graph attention networks," in Proc. Int. Conf. Learn. Representations (ICLR), 2018.

[17] P. A. Lopez et al., "Microscopic traffic simulation using SUMO," in Proc. IEEE Int. Conf. Intelligent Transportation Systems, 2018, pp. 2575–2582.

[18] A. Dosovitskiy et al., "CARLA: An open urban driving simulator," in Proc. 1st Annu. Conf. Robot Learning (CoRL), 2017, pp. 1–16.

[19] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," Phys. Rev. E, vol. 62, no. 2, pp. 1805–1824, 2000.