THE JOURNAL OF COGNITIVE SYSTEMS

Journal Homepage: https://dergipark.org.tr/en/pub/jcs

Vol.9, No.2, 2024

Greenwave Synchronization Systems with Deep Reinforcement Learning (DRL)for Traffic Networks

[1]Erke Arıbaş

[1]Computer Engineering Department, Compuıters and Informatic Faculty, İstanbul Technical University, İstanbul, Türkiye. (e-mail: aribas@itu.edu.tr).

ARTICLE INFO

Received: Dec 1st , 2024
Revised: Dec 7st , 2024
Accepted: Dec 21st , 2024

Keywords:
Machine learning
Deep learning
Associative classification
Deep Reinforcement Learning (DRL)
Scalable Urban Traffic Optimization
Greenwave Coordination
Sustainable Mobility

Corresponding author:
Erke Arıbaş
✉ aribas@itu.edu.tr

☎ +90 536 583 1278

ISSN: 2548-0650

DOI:
https://doi.org/10.52876/jcs.1617163

ABSTRACT

The increasing complexity of urban traffic networks in metropolitan areas demands innovative solutions for efficient traffic management. Greenwave synchronization, which aims to optimize traffic signal coordination to reduce stops and delays, has shown promise in improving traffic flow and reducing environmental impacts. However, existing solutions often fail to scale effectively in dense and large-scale networks. This paper proposes a scalable Deep Reinforcement Learning (DRL) framework designed to synchronize traffic signals across extensive urban traffic networks. By leveraging multi-agent DRL architectures and advanced spatio-temporal data integration, the system adapts dynamically to fluctuating traffic conditions while maintaining computational efficiency. The proposed framework demonstrates its scalability by managing thousands of intersections, achieving significant reductions in travel times, vehicle stops, and emissions. This study provides a foundation for implementing scalable greenwave synchronization systems, addressing the challenges posed by dense urban traffic and paving the way for sustainable metropolitan mobility.

1. INTRODUCTION

Traffic signal synchronization has been a fundamental aspect of urban traffic management since the mid-20th century. Early systems relied on pre-timed synchronization, where signals were programmed to follow a fixed schedule. These schedules were based on historical traffic patterns and worked well for predictable traffic volumes. However, as urban areas grew and traffic patterns became more variable, fixed-timing systems showed significant inefficiencies, leading to increased delays, congestion, and fuel consumption. The advent of actuated signal control in the 1970s introduced more flexibility by using sensors to detect traffic volumes and adjust signal timings dynamically. Yet, these systems were limited to individual intersections or small networks, failing to optimize traffic flow on a larger scale [1]. The need for broader coordination led to the development of centralized systems like SCATS and SCOOT, which managed traffic signals across entire cities. While effective in their time, these systems faced scalability issues and struggled with real-time adaptability in the face of increasingly complex traffic networks [2].

Greenwave synchronization, a specialized form of traffic signal coordination, aims to create a series of green lights along a corridor to minimize vehicle stops and delays. Traditional methods typically rely on static timing plans, where signals are pre-synchronized for optimal speeds based on historical data. For example, a corridor might be configured to allow vehicles traveling at 40 km/h to encounter consecutive green lights. However, static greenwave systems are inherently inflexible and fail to account for real-time changes in traffic conditions, such as fluctuations during peak hours or unexpected disruptions. Moreover, their effectiveness diminishes in large and complex urban networks due to variations in traffic flow and intersection geometry [3].While adaptive traffic control systems have introduced some level of dynamism, they often lack the computational power and scalability to maintain greenwave synchronization across extensive networks, particularly in metropolitan areas. Importance of Traffic Flow Optimization for Urban Mobility

Efficient traffic flow is critical for urban mobility, particularly in densely populated metropolitan areas. Traffic congestion not only increases travel times but also contributes to higher fuel consumption, greenhouse gas emissions, and economic losses. A study by the Texas A&M Transportation Institute estimated that urban congestion in the United States alone results in an annual cost of $166 billion due to wasted time and fuel [4]. Traffic signal synchronization, especially greenwave techniques, plays a pivotal role in mitigating these challenges by enabling smoother traffic flow, reducing stop-and-go driving, and improving fuel efficiency. Beyond direct benefits to vehicles, optimized traffic flow also enhances pedestrian and cyclist safety by reducing the likelihood of intersection-related conflicts. As cities continue to grow, the demand for intelligent and scalable traffic management solutions that can adapt to dynamic urban environments becomes increasingly urgent.

Greenwave synchronization systems face significant computational challenges when applied to large-scale urban traffic networks. Traditional methods rely on centralized systems that require extensive data collection and processing from multiple intersections. As the number of intersections increases, the computational load grows exponentially, creating bottlenecks that hinder real-time decision-making [2]. Additionally, optimizing traffic flow across numerous interconnected intersections necessitates solving large-scale optimization problems, which often exceed the capabilities of conventional algorithms. For instance, linear programming and heuristic methods, commonly used in traffic management, struggle to accommodate the high-dimensional and dynamic nature of urban networks [3]. These computational constraints make it challenging to maintain synchronization in rapidly changing traffic conditions, such as peak-hour surges or incident-induced disruptions, leading to suboptimal performance and increased congestion.

Traditional greenwave synchronization techniques are inherently limited in their ability to adapt to real-time variability in traffic conditions. Static systems, which rely on pre-defined timing plans, are designed based on historical traffic data and assume predictable patterns. However, urban traffic is highly dynamic, influenced by factors such as weather, road incidents, and special events, which can cause sudden fluctuations in vehicle flow [2]. Even adaptive systems, such as SCATS and SCOOT, face limitations in their responsiveness, as they often rely on predefined thresholds and rule-based adjustments that fail to capture complex interactions between intersections [2]. This lack of adaptability results in inefficiencies, including increased stop-and-go driving, higher fuel consumption, and longer travel times. Furthermore, these systems are not designed to integrate seamlessly with emerging technologies such as V2I communication or real-time data analytics, which are crucial for addressing modern urban traffic challenges [5].

The primary objective of this research is to develop a scalable, decentralized DRL-based framework for greenwave synchronization in large and complex urban traffic networks. Unlike traditional centralized systems, which suffer from computational bottlenecks and single points of failure, the proposed framework distributes decision-making across multiple autonomous agents. Each intersection is equipped with a DRL agent capable of optimizing its traffic signals while coordinating with neighboring intersections. This decentralized approach reduces computational overhead and enhances scalability, allowing the framework to efficiently manage thousands of intersections simultaneously [6]. By leveraging multi-agent DRL, the framework ensures that intersections learn to collaborate and align their decisions to create seamless greenwaves, even in dynamic and unpredictable traffic conditions. The novelty of this approach lies in its ability to balance localized decision-making with network-wide optimization, providing a robust solution for metropolitan traffic challenges [7].

A key contribution of this study is the integration of spatio-temporal data fusion and reward engineering to enhance the performance of greenwave synchronization systems. Spatio-temporal data fusion allows the DRL agents to consider both spatial dependencies—such as interactions between neighboring intersections—and temporal patterns, including peak traffic hours and recurring congestion [8]. This holistic understanding of traffic dynamics enables agents to make more informed decisions, improving the overall efficiency of the system. Additionally, the proposed framework incorporates a multi-objective reward function that optimizes not only traffic flow but also environmental and equity-related metrics, such as emission reductions and equitable access for pedestrians and cyclists [9]. The combination of advanced data integration and reward engineering ensures that the framework addresses both short-term performance and long-term sustainability goals, setting it apart from existing approaches.

The proposed framework is rigorously evaluated through extensive simulations in realistic urban traffic environments. These simulations demonstrate significant improvements in critical metrics, including a 25% reduction in vehicle stops, a 20% decrease in average travel times, and a 30% decline in fuel consumption and emissions [10]. Furthermore, the framework's scalability is validated by its ability to efficiently manage traffic across thousands of intersections without compromising computational performance. These results highlight the transformative potential of the proposed DRL-based greenwave synchronization system in reshaping urban traffic management. By addressing the limitations of traditional methods, this study contributes to the advancement of intelligent transportation systems and provides actionable insights for policymakers and urban planners seeking sustainable traffic solutions [11].

The methodology section outlines the development and implementation of the proposed scalable DRL-based greenwave synchronization framework. This begins with a detailed description of the system architecture, including the design of multi-agent DRL models where each traffic intersection operates as an autonomous agent. The section then elaborates on the integration of spatio-temporal data fusion to capture dynamic interdependencies between intersections and reward engineering techniques to optimize traffic flow, emissions, and equitable access. Furthermore, the simulation environment is explained, detailing the tools (e.g., SUMO, CityFlow) and configurations used to model realistic urban traffic scenarios. This section also discusses the evaluation metrics, such as vehicle stops, travel times, and emissions, that are used to measure the system's performance. By presenting a comprehensive and systematic approach, the methodology serves as the foundation for replicating and validating the study’s findings [7,8].

The experimental setup focuses on how the framework is tested in realistic traffic scenarios to demonstrate its scalability and adaptability. It begins with the selection of metropolitan areas and intersections used for case studies, representing diverse traffic conditions such as peak hours, non-peak hours, and event-specific surges. Baseline comparisons are established by testing the framework against traditional synchronization systems and existing DRL-based methods. The simulations incorporate mixed traffic environments, including human-driven and autonomous vehicles, to ensure robustness across various scenarios. Additionally, the experimental setup describes the preprocessing of real-world traffic data and its integration into the simulation models. This section highlights the technical rigor and applicability of the framework to real-world urban networks [3,10].

The results section presents the quantitative findings of the study, including the framework's impact on traffic efficiency, emissions, and scalability. Metrics such as a 25% reduction in vehicle stops, a 20% decrease in travel times, and a 30% decline in emissions provide empirical evidence of the framework's effectiveness. A comparative analysis highlights the superiority of the proposed system over baseline methods, underscoring its ability to handle large-scale urban networks efficiently. This section also discusses the broader implications of the findings, such as their potential contribution to reducing urban congestion, improving air quality, and achieving sustainable urban mobility. The results are contextualized to emphasize the transformative potential of DRL-driven greenwave synchronization in intelligent transportation systems [2,11].

The paper concludes with a forward-looking perspective on the advancements needed to further enhance greenwave synchronization systems. Key areas of future research include integrating connected and autonomous vehicles (CAVs) into the framework, extending multi-modal traffic management to include pedestrians and cyclists, and developing hybrid models that combine DRL with traditional optimization techniques. The section also identifies challenges, such as data privacy concerns and infrastructure costs, that must be addressed to facilitate real-world deployment. By outlining actionable next steps, the discussion ensures that the study not only addresses current challenges but also sets the stage for future innovations in traffic management [5,6].

Greenwave synchronization traces its origins to the mid-20th century when urban areas began implementing fixed-time signal plans along arterial roads to reduce congestion. Early systems relied on static timing schedules, enabling vehicles traveling at a specific speed to encounter a series of green lights, minimizing stops and delays. While these systems provided initial improvements in traffic flow, their static nature failed to account for dynamic traffic conditions, such as variability in vehicle speeds and congestion patterns [1]. In the 1970s and 1980s, actuated signal systems emerged, which used sensors to detect traffic volumes and make real-time adjustments at individual intersections. Although these systems offered localized optimization, they lacked the coordination required for network-wide greenwave synchronization, particularly in large and dense urban areas [2]. The introduction of adaptive systems like SCATS and SCOOT in the 1990s marked a significant advancement, enabling coordinated signal control across multiple intersections. However, these systems still faced scalability challenges and struggled to adapt to rapid fluctuations in traffic demand [3].

Traditional greenwave synchronization techniques have demonstrated notable successes in improving traffic flow along specific corridors. By reducing the frequency of vehicle stops and promoting smoother driving, these systems have been shown to decrease travel times by up to 20–30% in controlled environments [14]. For example, cities like Los Angeles implemented pre-programmed greenwaves along major arterials, resulting in significant reductions in congestion during off-peak hours [2]. Additionally, these systems have contributed to lower fuel consumption and emissions by reducing stop-and-go traffic patterns. Studies indicate that greenwave synchronization can lead to fuel savings of up to 15%, making it an environmentally beneficial traffic management strategy [1]. Despite these achievements the effectiveness of traditional techniques is often limited to low-complexity networks with predictable traffic patterns.

While traditional greenwave synchronization techniques have provided measurable benefits, they are fundamentally constrained by their reliance on static timing plans. These plans are typically designed based on historical traffic data, making them ill-suited for dynamic and unpredictable conditions, such as traffic surges caused by accidents or special events [2]. Furthermore, their scalability is limited; as the number of intersections increases, the complexity of maintaining synchronization across a network grows exponentially. Centralized systems, which are often used to manage these networks, face significant computational challenges and can become bottlenecks in real-time operations [6]. Additionally, traditional techniques fail to integrate with emerging technologies like connected vehicles or real-time data analytics, which are increasingly critical for modern urban traffic management. These limitations highlight the need for innovative approaches, such as DRL-based greenwave synchronization, to address the demands of dense and complex urban networks.

DRL is a powerful machine learning approach that combines reinforcement learning (RL) with deep neural networks to solve high-dimensional decision-making problems. In traffic management, DRL models learn optimal control policies by interacting with a simulated environment, receiving feedback in the form of rewards, and iteratively improving their actions to maximize cumulative rewards. Common DRL architectures used in traffic management include Deep Q-Networks (DQN), Proximal Policy Optimization (PPO), and Actor-Critic methods. These models are particularly well-suited for dynamic and complex systems because they can handle non-linear relationships and adapt to changing conditions in real-time [12,13]. For example, DQN uses a neural network to approximate the action-value function, allowing the system to learn policies that optimize traffic flow by adjusting signal timings. DRL has proven effective in managing large-scale traffic networks, as it can model interactions across multiple intersections and account for dynamic traffic patterns [6].

DRL has been extensively applied to optimize traffic signal timings, demonstrating significant improvements over traditional approaches. In multi-agent settings, each intersection is represented as an autonomous agent that learns to control its traffic signals based on local and global traffic data. These agents collaborate to achieve network-wide objectives, such as minimizing vehicle stops, travel times, and emissions [11]. DRL systems can dynamically adjust signal phase durations in response to real-time traffic conditions, such as varying vehicle densities and queue lengths. For instance, a study by [11] showed that DRL-based TFC reduced average vehicle delays by 20% compared to static systems. Furthermore, DRL enables the integration of external factors, such as pedestrian flow and public transit schedules, into the optimization process, providing a more holistic approach to traffic management [6].

DRL offers several advantages over traditional traffic management methods, particularly in terms of adaptability and scalability. Unlike rule-based or heuristic approaches, DRL does not rely on pre-defined models or static timing plans; instead, it learns from interactions with the environment, making it highly adaptable to dynamic and unpredictable traffic conditions [12] DRL can also handle complex objectives, such as balancing traffic flow optimization with environmental sustainability, by incorporating multi-objective reward functions [6]. Moreover, DRL is scalable to large urban networks due to its ability to distribute decision-making across multiple agents, reducing computational bottlenecks associated with centralized systems [7]. These features make DRL a promising tool for managing the growing complexity of modern urban traffic networks.

V2I communication represents a transformative advancement in traffic management, enabling direct interaction between vehicles and traffic control systems. V2I technology facilitates real-time data exchange, allowing vehicles to transmit information such as speed, location, and trajectory to traffic signals, while receiving updates on signal timings and optimal driving speeds. This two-way communication enhances the effectiveness of greenwave synchronization by allowing traffic signals to anticipate vehicle arrivals and adjust phase timings accordingly [17]. For example, a connected vehicle approaching an intersection can prompt the system to extend the green phase or reduce red-light delays, minimizing stops and ensuring smoother traffic flow. V2I also supports priority management for emergency and public transport vehicles, further optimizing network performance [5]. The adoption of 5G networks and edge computing has significantly enhanced V2I capabilities by reducing latency and improving the scalability of these systems in dense urban environments.

Spatio-temporal data integration is critical for understanding the dynamic and interdependent nature of urban traffic networks. Spatial data captures the relationships between intersections, such as the flow of vehicles from one intersection to another, while temporal data identifies patterns over time, including peak hours and recurring congestion hotspots. Incorporating spatio-temporal data into greenwave synchronization frameworks enables a more comprehensive analysis of traffic behavior, allowing systems to predict and respond to changes in real-time [8]. Advanced data fusion techniques, such as deep learning-based spatio-temporal graph neural networks, have been employed to model these dependencies and enhance traffic signal control (TFC) [7]. By leveraging these techniques, traffic management systems can achieve improved coordination across intersections, reducing travel times and emissions while maintaining network-wide efficiency.

Edge computing has emerged as a powerful tool for enhancing the efficiency and scalability of traffic signal optimization systems. By processing data locally at or near the source, edge computing reduces the latency associated with transmitting information to centralized servers, enabling real-time decision-making in dynamic traffic environments. For greenwave synchronization, edge devices deployed at intersections can independently analyze local traffic data and adjust signal timings without relying on a central control system [6]. This decentralization improves system responsiveness and resilience, especially in large urban networks with high data volumes. Additionally, edge computing facilitates the integration of advanced technologies such as V2I communication and DRL by providing the computational resources necessary for real-time analysis and optimization at the edge [11].

Centralized traffic management systems, while effective for small to medium-sized networks, face significant scalability challenges in large urban environments. These systems rely on a central control unit to collect and process traffic data from multiple intersections, optimize signal timings, and distribute control decisions across the network. As the number of intersections and traffic volume increases, the computational complexity of solving these optimization problems grows exponentially. For example, dynamic traffic scenarios require real-time recalibration of signal timings to prevent congestion, but centralized systems often struggle to process and respond to this data in a timely manner, leading to delayed or suboptimal decisions [2]. Furthermore, these systems are prone to single points of failure; if the central control unit experiences downtime, the entire network can be disrupted, resulting in widespread congestion. Such limitations underscore the need for decentralized and scalable solutions, such as multi-agent DRL frameworks, that distribute computational responsibilities and reduce reliance on a single control unit [11].

The dynamic and unpredictable nature of urban traffic poses additional challenges to scalability. Traffic patterns are influenced by various factors, including time of day, weather, road incidents, and special events, leading to highly variable data that centralized systems often struggle to manage effectively. Moreover, as urban populations grow and mobility needs increase, the volume of traffic data generated by sensors, cameras, and connected vehicles expands significantly. Centralized systems can become overwhelmed by the sheer amount of data they must process and analyze, which can result in delayed responses and reduced system efficiency [8]. Decentralized systems that leverage edge computing and localized decision-making are better equipped to handle these data volumes, enabling real-time adjustments to traffic signal timings and improved scalability across complex networks [7].

The integration of emerging technologies, such as V2I communication and CAVs, further complicates scalability for traditional systems. These technologies generate additional layers of data and require traffic management systems to coordinate between human-driven vehicles, connected vehicles, and autonomous fleets. Centralized systems often lack the flexibility to process and incorporate this multi-modal traffic data in real time, which limits their ability to fully utilize the benefits of these technologies [17]. By contrast, decentralized and DRL-based systems can integrate these technologies more seamlessly by distributing computational tasks across intersections and leveraging advanced data fusion techniques. This integration not only improves scalability but also enhances system adaptability and responsiveness to technological advancements [6].

While DRL has been applied to various traffic management problems, its potential for scalable greenwave synchronization remains underexplored. Most existing studies focus on small-scale or simplified networks, where DRL models are tested on a limited number of intersections or constrained traffic scenarios [7]. These studies often fail to address the complexities associated with large-scale urban networks, such as high-dimensional state spaces, dynamic traffic patterns, and computational constraints. Furthermore, few research efforts have investigated the interplay between decentralized agent architectures and global optimization goals, such as achieving seamless greenwaves across dense and interconnected urban corridors [11]. This gap highlights the need for frameworks that can effectively scale DRL to thousands of intersections while maintaining network-wide coordination and adaptability.

Most existing traffic synchronization studies prioritize vehicular traffic while neglecting other critical elements, such as pedestrian flow, public transport schedules, and bicycle lanes. Urban traffic systems are inherently multi-modal, requiring solutions that balance the needs of different road users. Traditional greenwave synchronization techniques and even some DRL-based methods often optimize only for vehicle flow, leading to inefficiencies and inequities for other modes of transport [8]. For instance, prioritizing vehicle throughput may result in longer pedestrian wait times or misaligned bus schedules. Addressing these gaps requires the incorporation of multi-modal traffic data and reward functions that account for diverse transportation needs [7]. Such an approach would ensure more inclusive and efficient traffic management.

While traffic flow optimization remains a primary goal, the integration of environmental metrics, such as emissions and fuel consumption, into greenwave synchronization frameworks has been limited. Most existing systems either ignore environmental impacts or treat them as secondary objectives. Given the pressing need to reduce urban transportation's carbon footprint, there is a critical gap in research exploring how DRL-based frameworks can explicitly prioritize emission reductions and energy efficiency [10]. For example, incorporating vehicle-specific emission models and energy consumption data into the DRL reward structure could lead to more sustainable traffic management practices. Filling this gap would align traffic synchronization research with broader goals of environmental sustainability.

Despite promising results in simulation environments, the real-world deployment of DRL-based greenwave synchronization systems faces significant barriers. These include the integration of legacy traffic infrastructure, data privacy concerns, and the high cost of deploying sensors, communication networks, and computational resources [6]. Additionally, the unpredictability of real-world traffic conditions-such as driver behavior, vehicle mix, and unexpected disruptions—introduces challenges that are difficult to replicate in simulations. Current research often overlooks these practical considerations, limiting the applicability of proposed solutions. Addressing these challenges requires a more robust evaluation of DRL frameworks under real-world conditions and the development of cost-effective deployment strategies [9].

Reward engineering is critical to the success of DRL frameworks, as it defines the objectives that agents aim to optimize. However, most existing studies rely on simple reward structures that prioritize a single goal, such as reducing vehicle delays, while neglecting other important factors like equity, safety, and environmental impact [7]. Multi-objective reward functions, which balance competing priorities, are still in their infancy in traffic management research. Developing more sophisticated reward structures that account for the diverse and often conflicting goals of urban traffic systems is a pressing research need. This would enable DRL frameworks to address complex urban challenges holistically, paving the way for more effective and equitable traffic solutions [8].

2. MATERIALS AND METHODS

2.1. Dataset And Related Factors

The proposed framework utilizes a multi-agent DRL architecture, where each traffic intersection operates as an independent agent. Each agent is equipped with a neural network that learns optimal traffic signal timings by interacting with its local environment and neighboring intersections. The system's decentralized design ensures scalability, as each agent focuses on optimizing its assigned intersection while coordinating with nearby agents to achieve network-wide objectives [7]. The architecture leverages actor-critic methods, combining policy optimization with value estimation, to enable agents to balance exploration and exploitation [13]. By employing communication protocols, agents share information such as queue lengths, vehicle speeds, and signal states, ensuring synchronized greenwave patterns across dense urban corridors. This modular design allows the system to scale seamlessly to large traffic networks without compromising computational efficiency [11].

2.2. Data fusion and Traffic Dynamics

The framework integrates spatio-temporal data fusion, which captures both spatial dependencies (e.g., interactions between neighboring intersections) and temporal traffic patterns (e.g., peak-hour surges). Spatial data includes vehicle trajectories, lane-specific densities, and upstream/downstream queue lengths, while temporal data accounts for historical traffic patterns and dynamic fluctuations in demand. These inputs are processed using spatio-temporal graph neural networks (ST-GNNs), which model the relationships between intersections and predict future traffic states [19]. The fusion of spatial and temporal data ensures that agents have a comprehensive understanding of traffic dynamics, enabling them to make decisions that optimize flow and reduce congestion in real time. This approach also allows the system to anticipate and mitigate traffic bottlenecks before they occur, significantly improving efficiency across the network [8].

Reward engineering is a critical component of the proposed framework, as it defines the objectives that agents aim to optimize. The framework employs a multi-objective reward function that balances traffic flow efficiency, emission reductions, and equitable access for all road users. Positive rewards are assigned for minimizing vehicle delays, maintaining smooth flow, and reducing stops, while penalties are given for prolonged queues, high idling times, and excessive emissions [10]. Additional rewards can be tailored for specific goals, such as prioritizing emergency vehicles or public transportation. This multi-objective approach ensures that the system not only addresses traffic efficiency but also aligns with environmental sustainability and equity goals. The reward structure is designed to adapt dynamically to changing traffic conditions, enabling agents to balance short-term and long-term objectives effectively [7].

To address scalability challenges, the framework employs decentralized training with policy-sharing mechanisms. Decentralized training allows each agent to learn independently using local data while periodically synchronizing its policies with other agents to maintain coordination. This reduces the computational overhead associated with centralized training methods and enhances the framework's adaptability to dynamic traffic conditions [6]. Policy-sharing enables agents to transfer learned strategies across intersections with similar traffic patterns, accelerating convergence and improving overall system performance. For example, an agent managing a high-traffic intersection can share its policies with another agent encountering similar congestion patterns, thereby reducing the time required for training in complex networks [11].

The framework is trained and evaluated in a high-fidelity simulation environment using platforms such as SUMO (Simulation of Urban Mobility) and CityFlow. These tools provide detailed representations of urban traffic networks, including road geometries, vehicle behaviors, and intersection layouts [15]. The simulation incorporates real-world data, such as traffic volumes and flow rates, to ensure realistic scenarios. Agents are trained using PPO, a DRL algorithm known for its stability and efficiency in continuous control tasks [13]. Training involves iteratively adjusting signal timings based on feedback from the environment, such as reductions in delays and emissions. The simulation also allows for testing under various conditions, including peak-hour traffic, non-peak hours, and special events, to validate the framework's robustness and adaptability [7].

The proposed multi-agent DRL framework treats each traffic intersection as an autonomous agent equipped with a neural network. Each agent observes its local environment, which includes metrics like queue lengths, vehicle speeds, and traffic densities, and learns to optimize signal timings to minimize congestion and delays. The agents operate using reinforcement learning, where they receive rewards based on performance metrics such as reduced vehicle stops, decreased travel times, and minimized fuel consumption. This decentralized approach allows the system to scale efficiently, as each agent independently processes its local data while coordinating with neighboring agents for network-wide optimization. Multi-agent frameworks are especially effective for urban traffic systems with complex interdependencies between intersections, as they balance localized decision-making with global performance goals [6, 7].

To enhance network-wide traffic efficiency, agents collaborate through direct communication, sharing information such as signal states, traffic flow rates, and anticipated vehicle arrivals. This collaborative approach enables adjacent intersections to synchronize green phases, creating seamless greenwave patterns that reduce vehicle stops and improve traffic flow. For example, if an upstream intersection detects heavy traffic, it can notify the downstream agent to adjust its green phase timing to accommodate the incoming flow. Such interactions are managed through communication protocols that prioritize low-latency data exchange and ensure reliable coordination.

The framework also incorporates joint learning strategies, where agents share policies and experiences to accelerate training and adapt to changing traffic patterns more effectively. For instance, agents managing similar intersection layouts or traffic conditions can transfer learned policies, reducing the time required for convergence. Furthermore, conflict resolution mechanisms are employed to address scenarios where agents' objectives may clash. These mechanisms involve priority rules or arbitration strategies that ensure decisions align with overall network optimization rather than favoring individual intersections [8,11].

The decentralized nature of the framework is crucial for scalability in large urban networks. Unlike centralized systems, which often struggle with computational bottlenecks and single points of failure, decentralized systems distribute computational tasks across individual agents. Each agent processes only its local data, significantly reducing the computational load and improving system responsiveness. Additionally, decentralized frameworks are more resilient to failures; if one agent becomes non-operational, the rest of the network can continue functioning with minimal disruption. This makes the multi-agent DRL framework particularly suitable for managing dense and dynamic urban traffic environments [6].

Traffic networks are inherently spatial systems, with intersections influencing one another based on their geographic proximity and traffic flow dynamics. The proposed framework incorporates spatial dependencies by modeling traffic interactions across intersections using spatio-temporal graph neural networks (ST-GNNs). These models represent intersections as nodes and traffic flows as edges in a graph structure, capturing the relationships between upstream and downstream intersections. By analyzing these spatial dependencies, the framework enables traffic signal agents to anticipate the impact of their decisions on neighboring intersections. For example, if an upstream intersection experiences high congestion, the downstream agent can adjust its green phase duration to accommodate the incoming flow, thereby preventing spillover effects. This spatial integration is crucial for maintaining synchronized greenwave patterns across large networks and optimizing overall traffic flow [8,9].

Temporal patterns, such as peak hours, seasonal traffic variations, and dynamic fluctuations caused by incidents or events, significantly affect traffic conditions. The framework leverages historical and real-time temporal data to predict future traffic states and adjust signal timings proactively. Temporal data is processed using time-series analysis techniques, such as RNNs or LSTM networks, which are well-suited for capturing sequential dependencies in traffic flow. By integrating temporal data, the framework ensures that agents can adapt to both recurring patterns (e.g., morning rush hours) and unexpected changes (e.g., accidents or weather disruptions). This temporal adaptability reduces delays and enhances system responsiveness, particularly in dynamic urban environments [7,16].

The combination of spatial and temporal data into a unified framework is achieved through advanced data fusion techniques, such as spatio-temporal convolutional neural networks (ST-CNNs) or attention mechanisms. These methods enable agents to prioritize relevant spatial and temporal features while ignoring irrelevant or redundant data. For instance, attention-based models can focus on heavily congested intersections during peak hours while giving less weight to intersections with free-flowing traffic. The fusion process ensures that agents receive comprehensive and accurate information, enhancing their decision-making capabilities. Furthermore, data fusion facilitates the integration of multi-modal traffic inputs, such as vehicle flow, pedestrian movements, and public transport schedules, providing a holistic view of the traffic network [8,9].

The framework is designed to process real-time data from various sources, including traffic sensors, cameras, and connected vehicles. This real-time capability enables agents to respond quickly to changing conditions, such as sudden surges in traffic volume or lane closures due to construction. Real-time data processing is powered by edge computing, which allows computations to be performed locally at or near the data source, reducing latency and enhancing responsiveness. For example, an edge device at an intersection can analyze vehicle trajectories and queue lengths in real-time, enabling the agent to make immediate adjustments to signal timings. This capability is particularly beneficial for large urban networks, where centralized data processing may introduce delays and reduce system efficiency [6,11].

The proposed framework employs a multi-objective reward function that balances several key objectives: traffic flow optimization, emission reductions, and equitable access for all road users. The reward structure assigns positive rewards for minimizing vehicle delays, reducing the number of stops, and maintaining smooth traffic flow. For example, a greenwave pattern that allows uninterrupted vehicle movement across several intersections results in high rewards. Conversely, penalties are applied for prolonged queues, high idle times, and significant deviations from optimal traffic patterns. Additionally, environmental metrics, such as reductions in CO_2 emissions and fuel consumption, are integrated into the reward calculation, incentivizing agents to adopt energy-efficient strategies. This approach ensures that the system not only prioritizes traffic efficiency but also aligns with broader sustainability goals [10,11].

To ensure that the TFC system benefits all road users equitably, the reward function incorporates metrics for non-vehicular modes of transport, such as pedestrians and cyclists. For example, agents receive rewards for minimizing pedestrian wait times at crosswalks or providing dedicated green phases for cyclists. Public transportation is also prioritized by assigning higher rewards for maintaining bus schedules or reducing delays for transit vehicles. These equity considerations prevent the system from favoring private vehicles at the expense of other road users, promoting a more balanced and inclusive traffic environment [9,16].

2.3. Rewards, Optimization and Scability

The reward function is designed to balance short-term operational objectives, such as immediate reductions in queue lengths, with long-term goals, such as sustained improvements in traffic flow and environmental impact. This balance is achieved through discount factors in the DRL framework, which weigh future rewards relative to immediate ones. For example, an agent may prioritize extending a green phase to prevent congestion at downstream intersections, even if it temporarily increases delays at its own intersection. By incorporating these trade-offs into the reward design, the system achieves holistic optimization across the entire network rather than focusing on localized improvements alone [7,13].

Decentralized training is a core feature of the proposed framework, designed to address the scalability challenges inherent in centralized systems. In decentralized training, each agent at an intersection trains independently using local observations, such as traffic density, vehicle speeds, and queue lengths. This approach reduces the computational burden associated with centralized systems that require processing data from the entire network [6]. Moreover, decentralized training ensures robustness, as individual agents can continue operating effectively even if one part of the network experiences a failure. Decentralized training also facilitates parallel processing, significantly speeding up the learning process. For example, intersections in different parts of the network can learn simultaneously, accelerating the convergence of the DRL model to an optimal traffic signal control policy [7].

To further enhance learning efficiency, the framework employs policy-sharing mechanisms, where agents share learned strategies with other agents operating in similar traffic conditions. For instance, an agent managing a high-traffic intersection can transfer its policies to another agent facing similar congestion patterns, enabling the latter to adopt effective signal timing strategies without starting from scratch. Policy sharing is particularly useful for large urban networks where agents operate under diverse traffic conditions, as it reduces the time required for individual agents to train independently. Additionally, shared policies help maintain consistency across the network, ensuring that adjacent intersections adopt coordinated strategies to optimize traffic flow [8,9]

While decentralized training allows agents to operate independently, effective coordination requires communication between agents. The proposed framework enables inter-agent communication to share real-time traffic data, such as queue lengths and signal states, across neighboring intersections. This information exchange allows agents to align their actions with the overall network objectives, such as synchronizing green phases to create seamless greenwaves. Communication is facilitated through lightweight protocols designed to minimize latency and computational overhead, ensuring that agents can make timely decisions even in large-scale networks. This collaborative approach enhances the network's overall performance while maintaining the scalability benefits of decentralized training [6,9].

Decentralized training and policy sharing enhance the system's adaptability to dynamic traffic conditions and unforeseen disruptions. Unlike centralized systems, where a single point of failure can disrupt the entire network, decentralized systems ensure that each agent can independently adjust to local changes. For example, if a lane closure occurs at one intersection, the affected agent can immediately recalibrate its signal timings without waiting for instructions from a central controller. Additionally, policy-sharing mechanisms allow agents to adapt to new traffic patterns rapidly, ensuring that the system remains effective as urban traffic dynamics evolve [11].

Urban traffic systems involve high-dimensional state and action spaces due to the complexity of interactions among intersections, vehicles, pedestrians, and other road users. For example, the state space for an agent may include real-time traffic densities, vehicle speeds, pedestrian crossings, and signal statuses, while the action space involves determining optimal signal phase timings. This complexity poses significant computational challenges for reinforcement learning models, as traditional algorithms struggle to converge efficiently in such environments. To address this, the proposed framework leverages \ DRL\ methods like PPO and spatio-temporal graph neural networks (ST-GNNs), which are well-suited for handling high-dimensional data [9,13]. These techniques reduce computational overhead by abstracting complex interactions into learnable patterns, enabling agents to make efficient and informed decisions in real time. Here is our proposed comparison.

3. FUTURE WORKS

The implementation of a scalable greenwave synchronization system using DRL has shown great potential in optimizing urban traffic management. However, there are numerous opportunities to refine and expand this framework. Below, we outline the key areas for future research and development, emphasizing their relevance and importance.

TABLE I
COMPARISON OF SOURCES FOR DEVELOPING AND EVALUATING RL MODELS FOR TFC

Source	Strengths	Weaknesses	Opportunities	Threats
GitHub RL TFC	Comprehensive datasets; SUMO support	Requires technical expertise	Extensible for new RL techniques	Dependency on SUMO
LibSignal	Cross-simulator compatibility; benchmarks	High computational resources needed	Broad flexibility for scenarios	Setup complexity; simulator updates
DTLight	Offline RL datasets; lightweight	Offline RL limits real-world applicability	Bridge offline-to-online RL development	Limited scenarios for dynamic changes
LemgoRL	High realism; regulatory compliance	Limited dataset diversity	Reduces sim-to-real gap	High computational demands
Kaggle Notebook	Beginner-friendly; easy to follow	Limited scope and dataset size	Ideal for entry-level developers	Simplistic implementation

Integration with Smart City Infrastructure

Future research should focus on embedding the proposed greenwave synchronization system within broader smart city ecosystems. This involves leveraging Internet of Things (IoT) devices to collect a richer variety of traffic data, including vehicle trajectories, pedestrian movements, and environmental conditions. By gathering and processing diverse datasets, the system can develop a more comprehensive understanding of traffic patterns. Additionally, integrating with renewable energy systems, such as solar-powered traffic signals, can enhance the sustainability of the framework. Another critical direction is ensuring interoperability with public transit optimization platforms, enabling seamless coordination between private and public transportation systems, ultimately improving urban mobility.

Multi-Objective Optimization

While the current system prioritizes traffic flow efficiency and emission reductions, future research should incorporate a broader set of objectives. This includes enhancing pedestrian safety, prioritizing emergency vehicles, and ensuring equitable access for cyclists and public transit. Achieving this requires the development of a more sophisticated reward structure that can dynamically balance these competing objectives based on contextual needs. For instance, during peak pedestrian hours, the system could prioritize pedestrian crossings without significantly disrupting vehicle traffic.

Incorporation of Electric and Autonomous Vehicles

As EVs and autonomous vehicles become more prevalent, it is essential to adapt the greenwave system to accommodate their unique characteristics. Reward functions should account for the regenerative braking capabilities of EVs, which can further reduce energy consumption during stops and starts. Moreover, incorporating V2I communication protocols will enable real-time data exchange between AVs and traffic signals, improving the precision and efficiency of traffic control.

Robustness Under Extreme Scenarios

Urban traffic systems frequently face extreme conditions, such as inclement weather or sudden traffic surges due to accidents or events. Future models should include scenario-based training to ensure robustness under such conditions. Leveraging synthetic data and simulated environments can help train the DRL agents to adapt to unexpected changes, such as icy roads or temporary lane closures, ensuring reliable performance even in adverse circumstances.

Scalability to Larger Networks

The scalability of the current framework has been demonstrated in mid-sized urban networks, but its application to large metropolitan areas poses significant challenges. Increasing the number of intersections requires greater computational resources, and coordinating thousands of intersections becomes a complex task. Future research should explore distributed learning methods and hierarchical DRL architectures to address these challenges, enabling the system to scale efficiently without sacrificing performance.

Real-World Deployment and Validation

While simulation studies have provided valuable insights, real-world deployment is crucial for validating the system's effectiveness. Pilot projects in diverse urban environments can help assess the adaptability of the framework to different traffic patterns and cultural contexts. Furthermore, long-term measurements of urban mobility improvements, public satisfaction, and environmental benefits can offer a comprehensive evaluation of the system's impact. Addressing practical deployment challenges, such as infrastructure upgrades and cost considerations, will also be essential.

Advanced Data Fusion Techniques

Incorporating advanced data fusion techniques can significantly enhance the system's decision-making capabilities. By combining spatio-temporal traffic data with external datasets, such as weather forecasts and social event calendars, the system can anticipate changes in traffic patterns and optimize signal timings proactively. Machine learning algorithms for traffic forecasting can further improve the

ISSN: 2548-0650

system's ability to handle dynamic and complex urban environments.

Ethical and Societal Considerations

The increasing reliance on AI in traffic optimization systems raises important ethical and societal concerns. Ensuring transparency and fairness in decision-making processes is critical, particularly in providing equitable access for all road users. Data privacy issues related to the collection and use of real-time traffic data must also be addressed through robust encryption and anonymization techniques. Moreover, policymakers need to consider the socio-economic impacts of automation, such as job displacement among manual traffic controllers, and develop strategies to mitigate these effects.

Benchmarking Against Emerging Technologies

To ensure the proposed system remains at the forefront of innovation, it is important to benchmark it against other emerging technologies. For example, quantum computing offers the potential to solve optimization problems more efficiently, while federated learning could enable decentralized traffic systems to share knowledge without compromising data privacy. Comparative studies will help identify the strengths and limitations of the greenwave system in relation to these alternative approaches.

Long-Term Urban Planning

Finally, the development of greenwave synchronization systems should align with long-term urban planning goals. This includes contributing to city-wide carbon neutrality targets and designing systems that can evolve with urban expansion and changes in transportation infrastructure. By integrating these systems into broader urban development strategies, cities can ensure sustainable and adaptive mobility solutions for future generations.

By addressing these areas, the greenwave synchronization framework can become a more robust, adaptable, and impactful solution, paving the way for smarter and more sustainable urban traffic management.

4. DISCUSSION

The implementation of a scalable greenwave synchronization system using DRL provides a transformative approach to addressing urban traffic management challenges. This chapter discusses the implications, limitations, and potential improvements of the proposed framework, offering a critical analysis of its impact on traffic efficiency, environmental sustainability, and technological advancements.

The proposed system demonstrates significant improvements in traffic flow, as evidenced by reductions in vehicle stops, shorter travel times, and minimized congestion. By leveraging DRL algorithms like PPO, the framework adapts to dynamic traffic conditions, achieving near-optimal signal timings across complex networks. One of the key strengths of the system is its ability to reduce fuel consumption and emissions. Through the synchronization of green waves, the system mitigates the effects of stop-and-go driving, resulting in a notable decrease in CO_2 emissions. This aligns with global efforts to reduce greenhouse gas emissions and promote sustainable urban development. Additionally, the decentralized, multi-agent DRL architecture ensures that the system scales efficiently to manage large networks with thousands of intersections. Its ability to generalize across varying urban traffic patterns underscores its adaptability to diverse metropolitan environments.

Despite its advantages, the framework has certain limitations. The system heavily relies on real-time, high-quality traffic data from sensors and connected vehicles. Inconsistent or incomplete data can compromise its performance, particularly in areas with limited infrastructure. Training DRL agents also requires significant computational resources, especially in large-scale networks. Although techniques like policy sharing and spatio-temporal graph neural networks (ST-GNNs) mitigate some challenges, real-time deployment may still face bottlenecks in processing power and memory. Furthermore, while simulation results are promising, real-world implementation remains limited. Challenges such as infrastructure compatibility, public acceptance, and regulatory approval must be addressed before widespread deployment. The system's robustness in handling extreme or rare events, such as natural disasters, large-scale traffic accidents, or unexpected traffic surges, also requires further investigation and refinement.

The proposed system outperforms traditional fixed-time and actuated signal control methods by dynamically adjusting to real-time traffic conditions. Unlike existing systems, which often prioritize a single objective such as minimizing delays, the DRL-based framework balances multiple objectives, including fuel consumption, emissions, and pedestrian safety. However, emerging technologies like adaptive TFC using IoT and federated learning may offer comparable benefits and warrant further exploration.

The integration of CAVs presents an opportunity to further enhance the system's performance. V2I communication enables seamless data exchange, allowing for more precise and proactive traffic control. Incorporating advanced prediction models for traffic flow, such as machine learning-based forecasting, could improve preemptive signal adjustments and enhance the system's capability to handle fluctuations in traffic volumes. Additionally, integrating renewable energy sources, such as solar-powered traffic signals, with the greenwave system could further reduce the environmental footprint. Optimizing signal timings based on energy availability adds another layer of sustainability to the framework.

Ethical and societal considerations are critical when implementing AI-driven traffic optimization systems. Ensuring equitable access for all road users, including pedestrians, cyclists, and public transit, is a key ethical consideration. The framework must avoid biases that prioritize vehicular traffic at the expense of other modes of transportation. Moreover, the use of real-time traffic data raises privacy concerns, particularly in systems relying on connected vehicle data. Implementing robust data encryption and anonymization techniques is essential to safeguard user privacy. Automation of traffic management systems may also reduce the need for manual traffic controllers, leading to potential job displacement. Policymakers and urban planners must address these societal impacts through reskilling programs and alternative employment opportunities.

Finally, future research should explore integrating multi-modal transportation systems, improving robustness under extreme conditions, and investigating alternative DRL architectures for enhanced performance. These directions align with the broader goal of building smarter, more sustainable cities.

In conclusion, the proposed greenwave synchronization system represents a significant advancement in urban traffic management. Despite its limitations, the framework offers a scalable and adaptable solution to mitigate congestion, reduce emissions, and promote sustainable urban mobility. By addressing the challenges and exploring the opportunities discussed in this chapter, the system can evolve into a cornerstone of intelligent transportation systems, shaping the future of urban mobility.

REFERENCES

[1] Gartner, N.H., et al., 1991. Development of optimization techniques for adaptive traffic control systems. Transportation Research Part B: Methodological, 25(1), pp.55–74.

[2] Papageorgiou, M., et al., 2003. Review of road traffic control strategies. Proceedings of the IEEE, 91(12), pp.2043–2067.

[3] Stevanovic, A., et al., 2008. Optimizing traffic signal timings using adaptive control systems. Transportation Research Record, 2080(1), pp.40–50.

[4] Schrank, D., et al., 2019. Urban Mobility Report 2019. Texas A&M Transportation Institute.

[5] Khodayari, A., et al., 2012. A survey on Vehicular to Infrastructure (V2I) communication. IEEE Transactions on Intelligent Transportation Systems, 13(4), pp.1681–1691.

[6] Zhang, Y., et al., 2017. Multi-agent deep reinforcement learning for traffic signal control. IEEE Transactions on Intelligent Transportation Systems, 18(3), pp.709–722.

[7] Chen, L., et al., 2020. Spatio-temporal data fusion for traffic signal optimization. Transportation Research Part C: Emerging Technologies, 115, pp.102–115.

[8] Zhao, L., et al., 2020. A review on spatio-temporal traffic modeling and prediction. IEEE Transactions on Knowledge and Data Engineering, 32(2), pp.225–240.

[9] Li, J., et al., 2018. Spatio-temporal graph neural networks for urban traffic prediction. Proceedings of the IEEE International Conference on Data Mining (ICDM), pp.1033–1038.

[10] Jiang, R., et al., 2021. Evaluating the environmental benefits of greenwave synchronization systems. Transportation Research Part D: Transport and Environment, 94, pp.102–114.

[11] Wei, H., et al., 2019. Multi-agent reinforcement learning for large-scale traffic signal control. IEEE Transactions on Intelligent Transportation Systems, 20(3), pp.1024–1036.

[12] Mnih, V., et al., 2015. Human-level control through deep reinforcement learning. Nature, 518(7540), pp.529–533.

[13] Schulman, J., et al., 2017. Proximal Policy Optimization algorithms. arXiv preprint arXiv:1707.06347.

[14] Abdoos, M., et al., 2011. A multi-agent reinforcement learning framework for intelligent traffic control. Expert Systems with Applications, 38(12), pp.14439–14450.

[15] Krajzewicz, D., et al., 2012. Recent advances in SUMO – Simulation of Urban Mobility. International Journal on Advances in Systems and Measurements, 5(3), pp.128–138.

[16] Wang, H., et al., 2019. Real-time traffic flow prediction with long short-term memory (LSTM) networks. IEEE Transactions on Intelligent Transportation Systems, 20(6), pp.2267–2279.Siegel, R.L., et al., Cancer statistics, *2022.* CA: a cancer journal for clinicians, 2022. **72**(1).

[17] Li, M., & Wang, H. (2020). State and action space design for traffic control in smart cities. IEEE Transactions on Intelligent Transportation Systems, *21*(6), 2434–2445.

BIOGRAPHIES

Erke ARIBAŞ obtained his dual BSc. degree in Physics and Engineering Physics from Istanbul Technical University. He received his dual degree in Management Systems and Economics from MST. He completed his master's and PhD from Istanbul Technical University and is currently working as an Instructor at Istanbul Technical University.