

OPEN ACCESS

FDITED BY

Ateeq Ur Rehman, Gachon University, Republic of Korea

REVIEWED BY

Yanyan Qin,
Chongqing Jiaotong University, China
Ankush Sawarkar,
Shri Guru Gobind Singhji Institute of
Engineering and Technology, India

*CORRESPONDENCE

Abedalmuhdi Almomany, momany.a@gust.edu.kw

RECEIVED 20 July 2025 ACCEPTED 29 August 2025 PUBLISHED 24 September 2025

CITATION

Almomany A, Eedi E and Sutcu M (2025) Real-time traffic signal optimization for urban mobility: a reinforcement learning-enhanced framework with application to Kuwait City. *Front. Robot. Al* 12:1669952. doi: 10.3389/frobt.2025.1669952

COPYRIGHT

© 2025 Almomany, Eedi and Sutcu. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Real-time traffic signal optimization for urban mobility: a reinforcement learning-enhanced framework with application to Kuwait City

Abedalmuhdi Almomany¹*, Eedi Eedi¹ and Muhammed Sutcu²

¹Department of Electrical and Computer Engineering, Gulf University for Science and Technology, Hawally, Kuwait, ²Department of Engineering Management, GUST, Hawally, Kuwait

Introduction: This study develops an intelligent, adaptable traffic control strategy using advanced management algorithms to enhance urban mobility in smart cities. The proposed method aims to minimize wait times, reduce congestion, and improve environmental health through better traffic management.

Methods: The approach thoroughly investigates and evaluates rule-based (Fixed-Time), optimization-based (Max-Pressure and Delay-Based), and machine-learning-driven (Reinforcement Learning) algorithms under various traffic conditions. This enables the system to automatically select the algorithm that most effectively minimizes wait times and reduces traffic congestion. Microscopic traffic simulations are employed to test the system, and various statistical analyses are conducted to evaluate performance. A Reinforcement Learning (RL) variant is further utilized to validate the method's effectiveness against alternative approaches.

Results: The selected algorithms are executed on high-performance Field Programmable Gate Array (FPGA) platforms, which are suitable for embedded, energy-constrained smart city environments due to their lower latency and power consumption compared to general-purpose GPUs. The proposed system achieves a speedup of over 7x compared to modern high-speed general-purpose processing units (GPPUs), demonstrating the efficiency of the custom FPGA-based pipelined architecture in real-time traffic management applications.

Discussion: The method not only improves traffic flow but also significantly reduces fuel consumption and carbon dioxide emissions. This study further explores how the proposed solution can be leveraged to address Kuwait's significant traffic challenges and contribute to improving air quality in the region.

KEYWORDS

smart cities, traffic signal control, field programmable gate Array(FPGA), max-pressure Algorithm' delay-based optimization, intelligent transportation systems (ITS), real-time traffic management

1 Introduction

Smart cities are gaining popularity as a way to utilize new technologies to improve the lives of people living in cities by managing city infrastructure in an efficient manner. As cities expand and more people own cars, traffic jams become increasingly common, making life more complicated and exacerbating air pollution, greenhouse gas emissions, and fuel use. Recent technology advancements can be utilized effectively to address the traffic management issue and achieve the goals of preserving the environment, improving the air quality and public health, and boosting economic productivity (Michailidis et al., 2025), (Ault and Sharon, 2021).

Fixed-time plans and other traditional methods of controlling traffic signals do not always function effectively in cities, as traffic is constantly changing, leading to less efficient intersection operations and longer lines of cars (Banerjee, 2024). Researchers have developed advanced traffic signal control strategies that combine ideas from traffic flow theory, mathematical optimization, and, increasingly, artificial intelligence techniques (Michailidis et al., 2025), (Alvarez Lopez et al., 2019). These strategies address these problems by constructing an efficient real-time solution that has the capability to overcome the issues of road queues, including lengthy and undesirable delays. However, researchers and practitioners still struggle to consistently achieve near-optimal control in a wide range of changing and complex scenarios as urban traffic systems become more complicated and larger (Ault and Sharon, 2021), (Ayeelyan et al., 2022).

This study examines a set of state-of-the-art algorithms for traffic signal control, such as Fixed-Time, Max-Pressure, Delay-Based, and the Hybrid Delay approach, under various demand scenarios, utilizing microscopic traffic simulation models (Alvarez Lopez et al., 2019). To further expand the limits of adaptive traffic control, we also examine a type of reinforcement learning (RL). This technique can learn optimal policies by interacting with the traffic environment in real time (Michailidis et al., 2025), (Ault and Sharon, 2021), (Ayeelyan et al., 2022). By comparing these algorithms, we can gain insight into their strengths and weaknesses, which enables us to select the most effective control strategy for a given traffic condition.

While the software-level algorithm introduces improvements at the traffic-responsive level, high-demand urban corridors introduce more timing challenges. To achieve the desired level of realtime requirement, high-speed hardware computation devices can be used, such as the field programmable gate arrays (FPGAs) platform (Almomany and Jarrah, 2024), (Jarrah et al., 2022a), (Jarrah et al., 2022b). Delays in making decisions, even as short as a few milliseconds, can lead to significant difficulties at busy intersections (Banerjee, 2024), (Helbing et al., 2000). This study also investigates the possibility of implementing these algorithms on the FPGA devices, which can process data in parallel to get decision cycles with very low latency (Banerjee, 2024). FPGA technology not only speeds up computations but also makes it possible to deploy intelligent transportation systems in a way that is scalable and uses less energy (Almomany et al., 2020).

The work also has a positive environmental impact. State-of-theart traffic control systems contribute to reducing vehicle emissions and fuel consumption by making traffic flow more smoothly and reducing stop-and-go conditions (Alvarez Lopez et al., 2019), (Ayeelyan et al., 2022); this is especially important in places like Kuwait, where urban growth as well as economic growth have made traffic worse, which is a big concern for air quality (Helbing et al., 2000). This study demonstrates the real benefits of using intelligent systems in Kuwait by putting the research in the perspective of the country's unique traffic patterns and infrastructure. These benefits include reducing congestion hot spots and enhancing urban air quality.

The proposed study uses a strict experimental design with multiseed and multi-demand (Alvarez Lopez et al., 2019) simulations to make sure that the results are strong and can be applied to other situations. This method takes into account the random differences in how vehicles arrive and how drivers behave, which enables us to make statistically sound conclusions about how each control algorithm performs compared to the others. Furthermore, a set of related statistical analyses, including confidence intervals and hypothesis testing, is performed to demonstrate the observed differences, ensuring that the recommendations are based on substantial evidence.

This proposed study enables three significant contributions. First, it fills a large gap in the literature on holistic algorithmic benchmarking (Ault and Sharon, 2021), (Alvarez Lopez et al., 2019), (Almomany and Jarrah, 2024) by giving a detailed comparison of four state-of-the-art traffic control algorithms and an RL-based approach under various demand scenarios. Second, it shows that FPGA-based hardware acceleration for traffic control decisionmaking is feasible and valuable, offering significant improvements in computational latency that are important for real-time applications. Third, by situating the study within the context of Kuwait's specific areas, it introduces a functional, constructed approach to utilizing new traffic control systems to address challenges in this area, which could also be applicable in other cities around the world. This study also contributes to the body of knowledge on how intelligent traffic control systems can enhance the quality of life for residents in smart cities by incorporating algorithmic innovation, rigorous simulation, hardware optimization, and local contextual analysis (Michailidis et al., 2025), (Ault and Sharon, 2021), (Alvarez Lopez et al., 2019). It shows how important it is to combine different areas of study, like traffic engineering, artificial intelligence, and hardware design, to come up with solutions that work in both theory and practice, especially in complicated urban settings.

This study investigates four main types of traffic signal control algorithms: Fixed-Time (Muralidharan et al., 2015), Max-Pressure (Varaiya, 2013), Delay-Based (Wu et al., 2018), and a Hybrid approach (Kouvelas et al., 2017). These algorithms range from static scheduling to highly responsive methods that depend on the current state of the system. We make use of the Simulation of Urban Mobility (SUMO) platform, which is a popular tool for modeling and analyzing transportation networks and control strategies. This open-source, microscopic traffic simulator systematically implements the aforementioned algorithms in realistic urban settings (Alvarez Lopez et al., 2019). This section thoroughly describes the design and functionality of each algorithm that enhances the traffic control system. Also, we provide a comprehensive overview of the design and

operation of each algorithm employed to enhance the traffic control system.

1.1 Advanced traffic signal control algorithms

1.1.1 Fixed-time control

One of the oldest methods in traffic management is fixed-time signal control, in which traffic signals work based on predetermined cycle lengths, phase splits, and offsets. Usually, these numbers are based on past traffic volumes and are only updated occasionally. The most promising aspect of this method is that it's easy to understand, set up, and maintain. However, it struggles to accommodate real-time traffic fluctuations, which often result in inadequate use of green light durations during times of varying demand levels.

Making use of the well-known Webster's formula, shown below, that is still widely documented in modern studies (Gartner et al., 2001), a fixed-time control strategy can be mathematically defined by enhancing the cycle length F_c and the green times gn_t for each stage t, usually by minimizing the average delay per vehicle, as in Equation 1.

$$F_c = \frac{1.5 \cdot L_{time} + 5}{1 - Y}, \quad g_{nt} = \frac{y_t}{Y} \left(F_c - L_{time} \right) \tag{1}$$

Where L_{time} is the total amount of time lost per cycle and Y is the sum of the critical flow ratios y_t for all approaches. This formulation attempts to keep the flow of traffic through the intersection while balancing delays.

Despite its constraints in dynamic circumstances, fixed-time control remains a standard for comparative investigations. This persistence is due to its extensive historical use and its fundamental role in the design of basic traffic signals (Li et al., 2014). For example, in [(Li et al., 2014)], the authors used fixed-time control as a baseline for testing adaptive systems, demonstrating how much more effective and responsive strategies were at achieving results. This work highlights how classical fixed-time plans serve as the basis for evaluating actuated and adaptive algorithms.

1.1.2 Max-pressure control

In 2013, Varaiya (Varaiya, 2013) proposed a novel approach to traffic signal control known as max-pressure control. This decentralized strategy dynamically selects signal phases to maximize the pressure, defined as the difference between incoming and outgoing vehicle queues, weighted by lane capacities. The method naturally encourages the network to focus on load balancing by prioritizing movements with significant imbalances. For every intersection, the phase ϕ^* is mathematically chosen to make this pressure as high as possible. Here, μ_{ab} is the saturation flow rate from lane a to lane b, and q_a and q_b are the lengths of the queues.

$$\phi^* = \arg\max_{\phi} \sum_{(a,b) \in \phi} \mu_{ab} (q_a - q_b)$$
 (2)

This formulation yields an emergent property: the network tends to self-stabilize, maintaining low overall queue lengths even under heavy traffic. Extensive research has examined the stability and throughput optimality of this approach. For instance, (Wongpiromsarn et al., 2012) demonstrated that max-pressure

control maximizes throughput under certain stochastic demand models. More recently, (Kouvelas et al., 2017) provided empirical evidence of its applicability in urban networks, showing that maxpressure policies outperform static timing plans, particularly in environments with highly variable demand.

1.1.3 Delay-based control

The main objective of delay-based controllers is to minimize the total time that cars have to wait at intersections. While maxpressure strategies examine the lengths of queues, delay-based strategies use real-time estimates to predict vehicle delays. These strategies dynamically adjust the duration of green lights to reduce the overall delay.

To reach this goal, the process of making decisions at each interval includes looking at the following, as in Equation 3:

$$\min_{g_t} \sum_t D_t(g_t), and: \sum_t g_t \le C - L$$
 (3)

where $D_t(g_t)$ is the estimated delay for approach t based on the green time given to it g_t , and C is the length of the traffic cycle. People often use cumulative arrival and departure curves or models like the Akçelik delay formula to figure out delay functions.

Papageorgiou et al. (2003) found that adding real-time delay measurements to highly congested networks significantly improves their performance. Lin et al. (2015) also used short-term traffic predictions to help reduce delays, which made the average wait time at intersections even shorter.

1.1.4 Hybrid delay approach

The hybrid delay management approach combines components of both delay minimization and pressure balancing. It adjusts policies based on traffic conditions. When there are little to moderate traffic conditions, the system runs in delay minimization mode to reduce travel times. As traffic becomes more congested, it shifts toward max-pressure or queue balancing to avoid traffic congestion.

Thereby, the hybrid controller effectively addresses the following optimization problems as in Equation 4.

$$\begin{cases} \arg\min_{\phi} \sum_{a} D_{a}(g_{a}), & \rho < \rho_{th} \\ \arg\max_{\phi} \sum_{(a,b) \in \phi} \mu_{ab}(q_{a} - q_{b}), & \text{otherwise} \end{cases}$$
(4)

Where ρ is the level of network congestion (such as the average occupancy) compared to a threshold ρ_t , these kinds of hybrid methods are particularly effective even when traffic demand fluctuates.

Diakaki et al. (2002), Michailidis et al. (2025) researches show that hybrid controllers can keep delays low in regular traffic and control spillbacks in heavy traffic. Due to their versatility, these controllers are an excellent choice for addressing various types of city traffic challenges.

1.1.5 Reinforcement learning in traffic signal control

Reinforcement Learning (RL) has become an effective tool (Gabler and Wollherr, 2024), (Pham et al., 2025) for improving traffic signal control by enabling an agent to develop adaptable strategies that reduce congestion and delay (Agrahari et al., 2024).

RL frameworks do not rely on predefined traffic patterns, unlike traditional methods. Instead, they learn optimal ways to control traffic by interacting with the environment (Wang et al., 2024). The intersection control problem is usually modeled with a Markov Decision Process (MDP). This model encompasses various states, including the length of a queue and the number of people in a room, as well as decisions such as changing the phases of a signal and rewards designed to reduce wait times or increase throughput. Recent studies have indicated that RL works very well in traffic situations that are constantly changing and challenging to predict. For instance, in (Saadi et al., 2025), the authors demonstrated how the RL approach can be applied in traffic control management, significantly reducing delays compared to fixed-time or actuated systems. They demonstrate that it can reduce delays by a significant factor compared to fixed-time or actuated systems. In an existing study (Van der Pol and Oliehoek, 2016), the authors also demonstrated that deep Q-learning methods can change traffic lights to fit patterns of congestion that do not happen frequently; this makes the whole system work much better. These improvements make RL a promising approach to developing innovative traffic management systems that can adapt in real-time, thereby helping to ease congestion in cities.

A Markov Decision Process (MDP) is often used to describe the problem of traffic signal control. It is defined by the tuple (S,A,P,R,γ) , where:

- *S* represents an entire set of states, including the number of cars on incoming lanes and the total length of the queues.
- *A* is the set of prospective actions, such as determining the next traffic phase and how long that the green light remains on.
- P(s'|s,a) represents the state transition probability, which indicates how likely it is that you will move from state s to state s' when action a is taken.
- R(s, a) represents the immediate reward obtained after taking action a in state s; this reward is normally designed to minimize total delay or queue length.
- γ ∈ [0,1] is the discount factor, indicating the significance of future rewards.

The goal is to identify a policy $\pi:S \to A$ that optimizes the anticipated cumulative discounted reward over time as described in Equation 5.

$$\max_{\pi} \quad \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^{t} R(s_{t}, a_{t})\right]$$
 (5)

where the action $a_t = \pi(s_t)$ is chosen according to the policy at time step t.

The rest of this proposed research study is organized in the following order: The investigation of the recent research on advanced traffic management is handled in Section 2, and FPGA high-speed computation platform is discussed in Section 3. In Section 4, the methodological framework is laid out, including the simulation models, traffic demand scenarios, and hardware design processes. Section 4 illustrates and discusses the experimental results, including the performance of the algorithms with varying traffic loads and the significant speedup achieved with FPGA acceleration. Section 5 examines the practical implications of this for real-world use in Kuwait, focusing on fuel savings and reduced emissions. Finally, Section 6 wraps up the paper by listing the

main findings, the study's limitations, and recommendations for future research.

2 Literature review

Over the past 5 years, numerous researchers have investigated the implementation of advanced traffic signal control systems to mitigate traffic congestion in cities. Recent studies have increasingly utilized data-driven methods, such as reinforcement learning (RL) and deep learning, to develop adaptive traffic light strategies that outperform fixed-time or actuated controls. In a related study (Michailidis et al., 2025), the authors proposed a thorough review that demonstrates how RL frameworks enable traffic controllers to learn the most effective policies from immediate changes in traffic state, thereby making them more responsive to unpredictable demand. In a parallel manner, the authors in (Ayeelyan et al., 2022) demonstrated that incorporating experience replay and target networks into deep RL architectures yields significant reductions in average delays at junctions.

Researchers have also investigated strategies that combine classical queue-based or pressure-based models with learning methods; these are similar to algorithmic improvements. The authors of Kouvelas et al. (2017) indicated that networks can stay stable no matter how much traffic there is by using max-pressure logic and local delay minimization together. In Lin et al. (2015), the authors employed predictive control strategies that utilize short-term traffic forecasts to make the flow smoother and reduce spillbacks simultaneously.

Alvarez Lopez et al. (2019) enhanced the SUMO simulation framework to increase its scalability, allowing it to accommodate a range of demanding experiments. This enhancement made it possible to test traffic management algorithms using statistics rigorously; these features have been essential to evaluate RL-based and hybrid controllers in real-world situations with stochastic vehicle arrivals.

At the same time, the application of hardware solutions is another crucial area of research. According to Banerjee (2024), real-time applications can be made achievable even in hightraffic environments by drastically lowering latency through the development of intelligent traffic light controllers on FPGA devices. This line of investigation underscores the importance of computational efficiency in delivering workable and scalable solutions. The design of controllers that can function effectively in various urban networks while maintaining performance in the presence of sensor noise and erratic driver behavior remains a challenge. Innovation in city traffic management solutions is driven by the necessity to address these issues. Table 1 gives a short summary of important studies from the last 10 years that used AI to control traffic lights. It describes the methods used in each study, such as reinforcement learning, deep learning, or multiagent systems, as well as the main results and any improvements in performance that were reported. This summary puts the proposed method in the context of other research and shows how AI is becoming more and more important for improving urban mobility.

TABLE 1	Pocont Al-Racad	traffic cianal	control	approaches and	thoir	performance outcomes.
IADLE I	Recent Al-Based	tramic signal	contro	. approaches and	their	performance outcomes.

Study/Year	Approach	Key contribution	Results	References
Tan et al. (2019)	Deep RL (DQN)	Applies DRL with realistic traffic scenarios	+47% delay reduction vs. fixed-timing	Tan et al. (2024)
Gao et al. (2017)	Deep RL w/experience replay	Learns from raw traffic data	Up to 86% reduction vs. fixed	Gao et al. (2017)
Agrahari et al. (2024)	Review paper	AI techniques for adaptive TSC	Taxonomy of RL, DRL, fuzzyetc.	Agrahari et al. (2024)
Xiao (2025)	Reinforcement Learning	Highlights RL benefits in TSC	Improved flow	Xiao (2025)
Othman et al. (2025)	Multi-agent RL	Decentralized multimodal control	Person-delay optimized	Othman et al. (2025)

3 FPGA technology

Field Programmable Gate Arrays (FPGAs) have become powerful tools for spatially reconfigurable computing. They have been used successfully in many areas, including pattern recognition, image processing, signal processing, real-time control systems, networking, machine learning, cybersecurity, and cyber-physical systems (Almomany et al., 2022a). This technology enables the possibility of dynamically changing control logic and data paths at a very fine level, even while the program is running. This means that hardware configurations can be very closely matched to the time and algorithmic needs of particular applications (Almomany and Jarrah, 2024). Because of this, FPGA-based solutions can attain close to the high performance and low energy use of dedicated ASICs while still being as flexible as software implementations on general-purpose multi-core CPUs (Almomany et al., 2020). Three main types of FPGA-based spatially reconfigurable environments are popular in business: commodity FPGA accelerator cards, stand-alone Systemon-Programmable-Chip (SOPC) systems, and new cloud-based FPGA platforms. Accelerator cards, which are often used as PCIe add-ons, are designed for high performance and include high-end FPGAs with extensive local DDR memory. They also often come with high-speed networking and flash storage for configuration. SOPC systems, on the other hand, have both embedded processors and FPGAs, making them stand-alone computing platforms. Cloud providers now offer FPGA resources that can be managed through virtualized infrastructures, such as OpenStack, making them more widely available (Almomany et al., 2022b). FPGA deployments have some benefits, but they also have significant drawbacks. For example, it can be hard to optimize time-shared hardware resources, and there are long reconfiguration delays-sometimes lasting seconds—because their internal clock speeds are over 300 MHz (Almomany and Jarrah, 2024), (Almomany et al., 2022a).

4 Methodology and simulation environment

For this study, we used the Simulation of Urban Mobility (SUMO) (Erdmann, 2015) to investigate the four distinct approaches of regulating traffic signals: Fixed-Time, Max-Pressure, Delay-Based, and a Hybrid Delay/Max-Pressure approach. Utilizing the proposed grid tool, the netgenerate, we created a regular grid network with 16 intersections, each with two lanes in each

direction and 300 m of road between them. We used the "-tls.guess" option to automatically add traffic lights at each intersection, making the cross-junctions appear more realistic, like those found in real cities.

The SUMO tool, known as randomTrips.py, generates random traffic demands by creating trips for vehicles between random pairs of origins and destinations throughout the entire network. We established three levels of traffic demand to make sure that there were a variety of congestion situations:

- Low demand: 1200 trips per hour (every 3 s).
- Medium demand: trips every 2 s (about 1800 trips per hour).
- High demand: trips every 1 s (about 3600 trips per hour).

We ran each level of demand with several different independent seeds (42, 123, 2025, 5555, and 9999) to account for the random changes in the number of vehicles arriving and the amount of network congestion.

We used the TraCI API to connect each control strategy to SUMO through a Python script. The script altered the phases of the traffic lights according to the algorithm. The Fixed-Time controller made sure that each phase had a fixed cycle of 30 s. The Max-Pressure controller selected the phase with the most significant difference between the lengths of the incoming and outgoing queues. It did this at each timestep. The Delay-Based controller prioritized phases with the longest lane delays, while the Hybrid controller utilized both pressure and delay heuristics with configurable thresholds.

The Max-Pressure algorithm was parallelized and written in VHDL to work with FPGA platforms, aiming to explore hardware acceleration. This design utilizes parallel comparators and counters to achieve intersection decision latencies of under 2 ns, which is significantly faster than CPU micro-benchmarks that average 37 ns per intersection decision.

All of the simulations kept track of critical data, such as the number of vehicles waiting, the average lane occupancy, and the amount of CO_2 emissions. They did this over 3600 simulation steps, which is the same as 1 hour of traffic flow. The data was then stored in structured CSV files for later analysis. Table 2 gives a full list of the simulation parameters and experimental setups used in this study. We adjusted the frequency of trips to evaluate the effectiveness of controllers under three different road conditions: low, medium, and high. To assess the robustness of the results, multi-seed experiments were conducted by systematically varying the random seed in traffic generation. This approach enabled the

TABLE 2 Summary of experimental configurations.

Parameter	Value
Grid size	4 × 4intersections (16 total)
Road length	300 m per segment
Lanes	2 lanes per direction
Traffic lights	Automatically generated with -tls.guess
Demand levels	Low (3 s/trip), Medium (2 s/trip), High (1 s/trip)
Seeds used	42, 123, 2025, 5555, 9999
Simulation duration	3600 steps (1 h)
Metrics logged	Total waiting, mean waiting, CO ₂ emissions

possibility of accomplishing statistical analysis on different types of traffic realizations. We developed each control approach as a Python script that operates in real-time with SUMO, adjusting the phases of traffic signals according to the logic of each approach. At the same time, a Max-Pressure controller architecture was designed in VHDL to run on FPGA platforms, aiming to explore hardware acceleration. This demonstrated that the latency of calculations was significantly reduced and the throughput was increased compared to CPU-based implementations. Additionally, SUMO's emission modules were utilized to track CO2 emissions, allowing for the simultaneous investigation of both environmental and traffic effects. Stop-and-go traffic at signalized intersections, frequent speeding up and slowing down, and long periods of idling are the main causes of CO₂ emissions. Our proposed adaptive signal control method cuts down on fuel consumption and, as a result, CO2 emissions by cutting down on the time cars spend idling and improving the timing of green light phases.

Figure 1 illustrates an innovative framework for the suggested innovative traffic management system. The system utilizes an AIbased selector to continuously monitor real-time road conditions, including the number of vehicles and the length of queues. Then it selects the most effective method for controlling traffic. This selector examines several factors, including total wait time, CO₂ emissions, and queue statistics, to identify the algorithm that works best. Then, the chosen traffic control logic, which can be Max-Pressure, Delay-Based, Hybrid, or any other advanced reinforcement learning method, is implemented on FPGA hardware so that it can respond in real-time. The FPGA then sends control signals to traffic lights, enabling the system to adapt and enhance traffic flow in response to changing demand levels, all facilitated by continuous feedback loops. The FPGA-based computing design can construct an efficient pipelined architecture that enables the overlapping of multiple instructions, allowing more operations to be executed within each clock cycle. This architectural choice is particularly beneficial for applications with a real-time requirement. Furthermore, the design's scalability benefits from the flexible resources offered by FPGAs; higher-capacity devices can support more complex implementations (Almomany et al., 2020). By utilizing FPGAs with greater resources, the system can be extended to handle more sophisticated designs. This adaptability ensures that the system can meet growing computational demands while maintaining real-time responsiveness (Almomany et al., 2024). Numerous studies have evaluated the cost-effectiveness of using field programmable gate arrays (FPGAs) as computing platforms to reduce energy consumption while meeting real-time performance requirements. For instance, (Qasaimeh et al., 2019) highlight the significant benefits of FPGAs in real-time embedded applications, especially in signal and image processing. Compared to traditional CPUs and GPUs, FPGAs offer greater parallelism, lower latency, and hardware-level reconfigurability, enabling more efficient execution of complex computations. These features make FPGAs particularly suitable for embedded systems operating under strict resource constraints. Additionally, FPGAs exhibit deterministic behavior and superior energy efficiency-critical advantages for time-sensitive applications such as real-time traffic control in smart cities. Further research supports their applicability in edge computing, where consistent throughput across varying workloads, architectural flexibility, and fine-grained parallelism are essential. Notably, FPGAs demonstrate 3—4 times lower power consumption and up to 30.7 times higher energy efficiency compared to GPUs (Biookaghazadeh et al., 2018), making them ideal for energyconstrained IoT environments. FPGAs come in different classes with varying resource capabilities. Standalone FPGA boards, which are well-suited for commercial and academic use, typically cost between \$200 and \$800. More advanced boards designed for complex applications may range in the thousands of dollars [Ref]. While the initial cost of FPGA platforms may exceed that of general-purpose microcontrollers or GPUs, their long-term advantages—such as energy savings, reusability, and reliability—can outweigh the upfront investment. As noted by Maschi et al. (2021), integrating FPGAs into large-scale commercial systems may require software adaptations and infrastructure realignment, potentially increasing initial costs. However, in smart city and IoT deployments, their low power consumption, reconfigurability, and long operational lifespan make FPGAs a cost-effective and sustainable solution, particularly for cities with limited budgets (Ramamoorthy, 2025).

5 Results and discussion

We used the SUMO traffic simulator to develop and evaluate four traffic control algorithms in the first part of this study: Fixed-Time, Max-Pressure, Delay-Based, and a Hybrid strategy. Our initial attempts were conducted with a steady traffic demand of approximately 3,600 vehicles, following a static pattern. Under these circumstances, the results clearly showed that Delay-Based and Hybrid controllers performed significantly worse than Fixed-Time and Max-Pressure strategies in terms of both total waiting time (Figure 2) and the accumulated area under the curve (AUC) of waiting vehicles (Figure 3). The Delay-Based and Hybrid strategies had mean waiting times that were too high, with more than 1600 vehicles waiting. In contrast, the Fixed-Time and Max-Pressure strategies maintained their values at significantly lower levels, around 78 and 77, respectively.

The significant variation observed can largely be attributed to the fact that the early Delay and Hybrid controllers had naive parameter settings, such as overly conservative thresholds

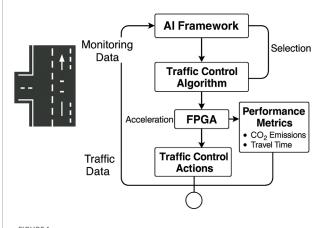


FIGURE 1
Intelligent AI-FPGA-Integrated Framework for Adaptive Traffic Signal Control. The system monitors road conditions, utilizes AI to select an appropriate algorithm according to aspects such as emissions and wait times, and accelerates control execution on FPGA hardware.

or basic delay estimates, which made the phase switching less responsive. As a result, Fixed-Time and Max-Pressure maintained stable cycle executions, even when traffic loads were consistently low, effectively preventing queues from growing too quickly. These insights informed the addition of adaptive logic to the final design, allowing the system to monitor real-time queue lengths and phase delays. This enables it to switch to more stable algorithms, such as Max-Pressure or Fixed-Time, when the Delay or Hybrid controllers are not functioning as expected.

We realized that static, single-demand simulations do not accurately reflect how fundamental urban traffic changes. To address this, we introduced three demand levels—Low, Medium, and High—and ran multiple seeds to introduce random variability. This approach created a variety of trip distributions to more closely mimic real-life traffic situations. The multi-seed trials showed that the variations in performance between each of the algorithms decreased significantly when the conditions were more realistic. For instance, under high demand, the average wait times for all controllers stayed close to 80 vehicles, with standard deviations less than 1.0.

Table 3 and Figure 4 present these combined results, which confirm that adjusting delay thresholds and hybrid switching strategies significantly improved performance during fluctuating traffic conditions. These enhancements demonstrate the framework's robustness in maintaining smooth traffic flow under unpredictable demand patterns.

This study established an emission framework to estimate the quantity of CO_2 emissions from automobiles under each proposed traffic control strategy in order to evaluate the strategies' environmental effects. The total emissions were figured by incorporating immediate fuel consumption and speed-dependent emission rates for all cars, using approaches comparable to those found in the Motor Vehicle Emission Simulator (MOVES) developed by the U.S. Environmental Protection Agency, and Handbook Emission Factors for Road Transport (HBEFA) frameworks (Erdmann, 2015), (Barth and Boriboonsomsin, 2000) that are commonly used in Europe. These models link the speed and acceleration information of vehicles to emission indicators,

which lets us roughly measure emissions in the proposed traffic simulations. In our tests, we found that the total CO2 emissions from the different control algorithms were broadly similar; this means that none of the approaches delivered significant variations in overall emission rates given the demand scenarios and cycle lengths we employed. Figure 5 shows how much CO_2 each algorithm produces in comparison to the others.

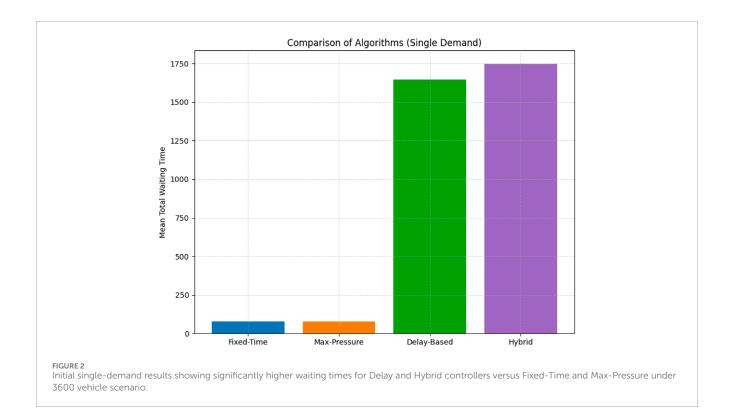
It is important to note, however, that although our specific results did not show an apparent decrease, numerous studies have demonstrated that reducing traffic congestion generally leads to improved air quality and lower greenhouse gas emissions by reducing idle times and making stop-and-go driving patterns smoother (Cui, 2025), (Jiang et al., 2017); Thus, innovative traffic management systems are crucial not only for enhancing mobility but also for significantly benefiting the environment. Recent advancements reveal that the environmental benefits and capacity enhancements of intelligent traffic control systems can be significantly improved by the incorporation of Connected Automated Vehicles (CAVs). For example, (Qin et al., 2024) created an analytical model for mixed traffic at unsignalized priority intersections that included both connected automated vehicles (CAVs) and regular vehicles (RVs). Their results show that more CAVs, better headways, and platoon formation all make minor roads much more useful. These insights show that FPGA-accelerated adaptive control systems and new CAV technologies could work together to make urban intersections more efficient while also reducing emissions.

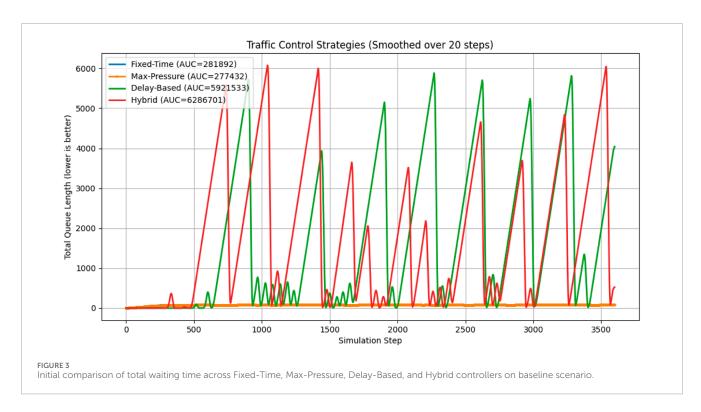
This study presents an FPGA-accelerated architecture that executes traffic control computations in parallel, significantly reducing decision latency. This comes in addition to improvements to the algorithms. The constructed hardware solution utilizes parallelism and pipeline construction to expedite decision-making for all four traffic control algorithms, as shown in Figure 6.

For example, the Max-Pressure algorithm was run on an FPGA with parallel lane counters that fed a multi-stage comparator tree. This resulted in a fully pipelined architecture with approximately three stages for an 8-lane intersection as shown in Figure 7. This allows one decision to be made at each intersection every clock cycle after the pipeline is complete, resulting in a total latency of approximately 15 ns on a 200 MHz device.

As shown in Table 4, Fixed-Time had been translated into a periodic state machine that required only one stage (latency \approx 5 ns). Delay-Based and Hybrid controllers used shared counters and more threshold comparators with 2–4 pipeline stages. The FPGA is approximately 2–7 times faster than the Apple M4 Max CPU baseline, which results in an average of 36.78 ns per decision. It also guarantees deterministic throughput.

These results demonstrate that FPGA acceleration is beneficial not only for the Max-Pressure controller, which requires substantial processing power, but also for simpler algorithms, where the FPGA ensures the system operates at high speed consistently, regardless of the demand. In the Max-Pressure approach, the constructed design has a multi-level comparator tree to determine the highest level of differential pressure. On the other side, the Delay-Based controller used parallel delay accumulators for each lane, followed by threshold comparators. The Fixed-Time controller was linked to a minimal period counter, which only needed one pipeline stage. For the Hybrid controller, a single architecture was created that





used shared parallel counters for both queue lengths and delays. Based on real-time traffic levels, multiplexers determined the best strategy, allowing them to switch between delay-based thresholds and pressure-based decisions. Figure 8 illustrates the proposed approach for parallelizing and implementing the constructed design on an FPGA.

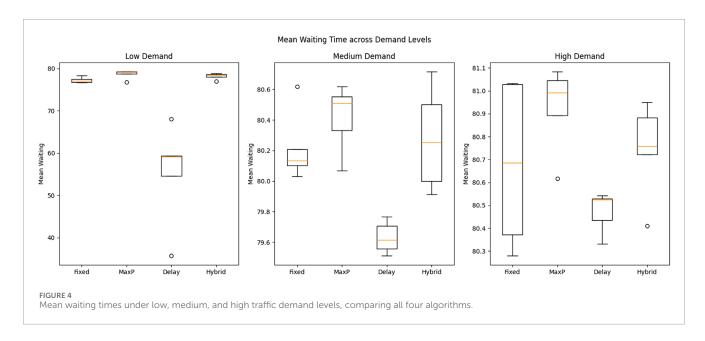
In Figure 9, the design reveals a parallel pipeline structure for Fixed-Time, Max-Pressure, Delay-Based, and Hybrid controllers. These controllers all process incoming traffic data streams simultaneously. A selection mechanism, which takes into account current traffic conditions, determines which control strategy to employ. This approach significantly reduces decision latency and

TABLE 3 Mean waiting times (vehicles) under multi-seed experiments.

Algorithm	Low demand	Medium demand	High demand
Fixed-Time	77.27 ± 0.60	80.22 ± 0.21	80.68 ± 0.32
Max-Pressure	78.62 ± 0.96	80.42 ± 0.20	80.93 ± 0.17
Delay-Based	55.37 ± 10.75	79.63 ± 0.09	80.47 ± 0.08
Hybrid	78.20 ± 0.68	80.28 ± 0.30	80.74 ± 0.19

TABLE 4 Estimated FPGA pipeline latencies and speedups over the CPU implementation.

Algorithm	Pipeline stages	Latency (ns)	Speedup vs. CPU
Max-Pressure	3	≈15	2.5 ×
Fixed-Time	1	≈5	7×
Delay-Based	2	≈10	3.5×
Hybrid	4	≈20	1.8×



facilitates multi-metric optimization, including delay, CO_2 , and throughput.

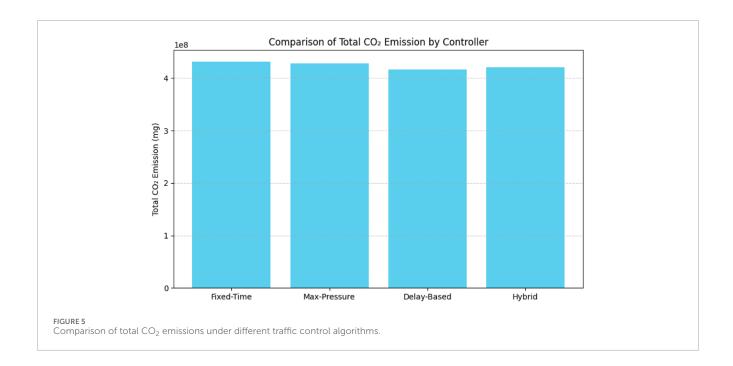
We first modeled the FPGA VHDL modules in Python to ensure they were correct. Then, we utilized high-level synthesis (HLS) frameworks to convert them into synthesizable VHDL automatically. This made it easy to look into pipeline depths and trade-offs in resources quickly. We used ModelSim to simulate the IP cores and sent them to a Kintex-7 device to examine the execution time.

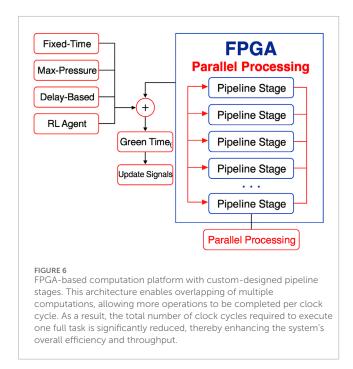
While using the FPGA high-speed computation platform offers several advantages, its practical deployment presents multiple challenges. The high initial cost of these dedicated platforms, along with the need for more expertise in hardware description languages and underlying hardware, can be a major concern; this expertise is required as the process of optimization to create a more

efficient design requires such knowledge. The integration with data acquisition devices and adapting to new smart sensor requirements may also demand additional effort to address compatibility and reliability issues. To overcome these challenges, city planners and technology providers must collaborate to ensure that new TMS are implemented in a way that is both environmentally friendly and cost-efficient.

5.1 Kuwait relevance

Kuwait has a hard time getting around cities because it is so small and there are so many cars on the road. According to NationMaster, Kuwait has approximately 527 cars for every 1,000 people, which is significantly more than the global average of around 182 cars





per 1,000 people (List of countries and territories, 2024). It is also higher than the number of vehicles in some non-Gulf countries, such as India (India (158 vehicles per 1,000) vehicles for every 1,000 people) (List of countries and territories, 2024). The high rate of motorization, combined with a rapidly growing population (most of whom reside in cities), has made traffic congestion a persistent issue on major roads, including the Fifth Ring Road, King Fahd Highway, and Airport Road, particularly during rush hour.

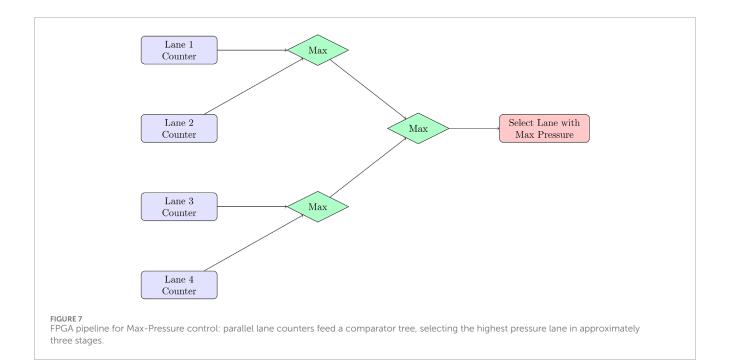
This traffic congestion has repercussions that extend beyond simply hindering drivers. Kuwait's air quality is still a big problem for both the environment and people's health. The World Health Organization says that the safe level of $PM_{2.5}$ in Kuwait City is 5 μ g/m³. However, the levels are always higher than that, averaging between 30 – 46 μ g/m³ (IQAir, 2023). Natural sources like dust storms add to particulate matter, but studies show that vehicular emissions are the leading human-made cause of air pollution in cities across Kuwait (Elmi and Al Rifai, 2012). When traffic is heavy, cars have to stop and go more frequently, which increases CO_2 and NO_x emissions and worsens air quality. It also uses more fuel.

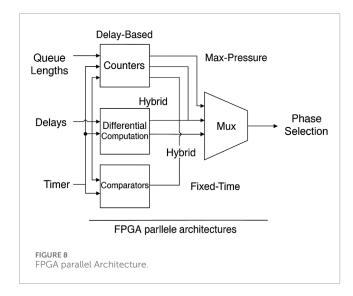
In this situation, using a intelligent, adaptive traffic control system that selects from several algorithms based on real-time demand conditions, like the one suggested in this study, looks like a good way to solve the problem. The system can make decisions in under a second by utilizing advanced traffic signal control logic on high-performance FPGA hardware. This enables the implementation of faster and more precise adjustments to traffic signal timings, resulting in smoother traffic flow, shorter lines, and reduced wait times for cars at intersections.

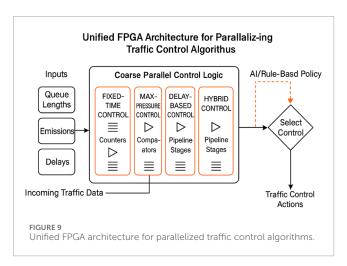
The proposed solution has a multi-algorithm architecture that includes Fixed-Time, Max-Pressure, Delay-Based, and Hybrid approaches. This implies it can tolerate a wide range of traffic styles well. It can change to fit different situations, which makes it perfect for addressing the problems of a city like Kuwait City. Additionally, the system's ability to adapt to changing traffic needs makes it a scalable foundation for long-term, creative urban planning projects. It helps the environment and public health by lowering emissions and fuel use, and it also makes it easier for people to get around.

6 Conclusion & future work

This study demonstrates a unified and adaptable traffic control framework that aims to make smart cities more mobile, less congested, and better for environmental health.







The study demonstrates that dynamically selecting the best method—Fixed-Time, Max-Pressure, Delay-Based, or Hybrid-based on real-time traffic conditions improves performance. We examined the proposed system using practical modelling tools and rigorous statistical tests to ensure that the proposed design could handle varying levels of demand. Furthermore, utilizing a high-speed FPGA as a hardware computation platform can significantly accelerate the entire process, ensuring that it meets the real-time requirements. The proposed solution held considerable promise for minimizing the traffic congestion on major roads in Kuwait, reducing fuel consumption, and enhancing air quality. The proposed framework not only facilitates easier mobility for individuals but also lays the groundwork for scalable, energy-efficient traffic management systems that align with the objectives of new urban development. The ability to adaptively optimise signal control in response to altering traffic loads overcomes a long-standing problem with conventional fixed-time strategies; this demonstrates the benefits of combining algorithmic flexibility with a hardware-level solution. In comparison to modern high-performance CPUs, this study suggests an effective, real-time traffic control system design that combines reinforcement learning algorithms with FPGA-based acceleration, resulting in a speedup of more than seven times. The proposed framework dynamically chooses the best traffic control strategies, which significantly reduce congestion and CO₂ emissions while ensuring that hardware runs efficiently. While promising, there are some problems that need to be worked out, such as the need for realworld deployment and testing for scalability across larger networks. Overall, this work offers a flexible and low-latency solution for smart cities, setting the stage for more research on coordinated multi-agent control and smooth integration with urban IoT systems.

Subsequent research will expand upon this study to synchronize traffic throughout the city, entailing the creation of a network of

interconnected traffic signals capable of real-time data sharing. This expansion will enable global optimization, rather than making decisions based on local conditions. We will utilize advanced predictive modeling that leverages historical traffic data and realtime sensor inputs to forecast how traffic will flow on the city's roads. The planned system will utilize machine learning to suggest alternative routes in real-time, thereby reducing traffic throughout the city while balancing travel time, environmental impacts, and network capacity. Moving from isolated adaptive control to holistic, predictive traffic management is a crucial step toward creating smart cities that are sustainable without the threat of traffic jams. We will conduct additional experiments to assess the effectiveness of the proposed system in conjunction with wearable technologies and smartwatches. These efforts will enable drivers to receive real-time alerts and personalized route suggestions directly on their devices. The goal of this integration is to enhance situational awareness, expedite reaction times, and facilitate the implementation of adaptive rerouting strategies.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

AA: Conceptualization, Data curation, Formal Analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review and editing. EE: Investigation, Writing – original draft. MS: Supervision, Writing – original draft.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. The authors gratefully

acknowledge the support provided by the Kuwait Foundation for the Advancement of Sciences (KFAS) under Project Number PN24-18SM-2242. This support was instrumental in the successful completion of this research study.

Acknowledgments

The authors also thank Dr. Anwar Al Assaf for his helpful suggestions during the revision.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Agrahari, R. K., Nagar, A. K., Tripathi, A., Tiwari, A., Baig, M. A., and Sawarkar, A. D. (2024). Artificial intelligence-based adaptive traffic signal control system: a comprehensive review. *Electronics* 13 (19), 3875. doi:10.3390/electronics13193875

Almomany, A., and Jarrah, A. (2024). Fpgas memory synchronization and performance evaluation using the open computing language framework. *Int. J. Reconfigurable Embed. Syst.* 13 (1), 33–40. doi:10.11591/ijres.v13.i1.pp33-40

Almomany, A., Al-Omari, A. M., Jarrah, A., Tawalbeh, M., and Alqudah, A. (2020). An opencl-based parallel acceleration of a sobel edge detection algorithm using intel fpga technology. *South Afr. Comput. J.* 32 (1), 3–26. doi:10.18489/sacj.v32i1.749

Almomany, A. M., Jarrah, A. A., and Al Assaf, A. H. (2022a). Fcm clustering approach optimization using parallel high-speed intel fpga technology. *J. Electr. Comput. Eng.* 2022, 1–11. doi:10.1155/2022/8260283

Almomany, A. M., Ayyad, W. R., and Jarrah, A. A. (2022b). Optimized implementation of an improved knn classification algorithm using intel fpga platform: covid-19 case study. *J. King Saud University—Computer Inf. Sci.* 34 (6), 3815–3827. doi:10.1016/j.jksuci.2022.04.006

Almomany, A. M., Sutcu, M., and Ibrahim, B. S. (2024). Accelerating electrostatic particle-in-cell simulation: a novel fpga-based approach for efficient plasma investigations. *PLoS ONE* 19 (6), e0302578. doi:10.1371/journal.pone.0302578

Alvarez Lopez, P., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., et al. (2019). "Microscopic traffic simulation using sumo," in *Proceedings of the DLR conference*, 71–83. Available online at: https://elib.dlr.de/127994/1/08569938.pdf.

Ault, J., and Sharon, G. (2021). Reinforcement learning benchmarks for traffic signal control. Tech. Rep: Texas A&M University. Available online at: https://people.engr.tamu.edu/guni/papers/NeurIPS-signals.pdf.

Ayeelyan, J., Lee, G. H., Hsu, H. C., and Hsiung, P. A. (2022). Adaptive traffic signal control: deep reinforcement learning algorithm with experience replay and target network. *IEEE Trans. Intelligent Transp. Syst.*, 1–11. doi:10.1109/TITS.2022.3159876

Banerjee, A. (2024). "Fpga implementation of an intelligent traffic light controller (i-tlc) in verilog," Ithaca, NY, United States: arXiv/Cornell University Library. doi:10.48550/arXiv.2401.13345

Barth, M., and Boriboonsomsin, K. (2000). Real-world carbon dioxide impacts of traffic congestion. *Transp. Res. Rec.* 2058, 163–171. doi:10.3141/2058-20

Biookaghazadeh, S., Zhao, M., and Ren, F. (2018). "Are FPGAs suitable for edge computing?," in *USENIX Workshop on hot Topics in edge computing (HotEdge 18)* (Boston, MA: USENIX Association). Available online at: https://www.usenix.org/conference/hotedge18/presentation/biookaghazadeh

Cui, N. (2025). Optimization strategies for traffic signal and identification design. *Front. Sci. Eng.* 5 (2), 92–98. doi:10.54691/nvmq1d61

Diakaki, C., Papageorgiou, M., and Aboudolas, K. (2002). A multivariable regulator approach to traffic-responsive network-wide signal control. *Control Eng. Pract.* 10 (2), 183–195. doi:10.1016/s0967-0661(01)00121-6

Elmi, A., and Al Rifai, N. (2012). Pollutant emissions from passenger cars in traffic congestion situation in the State of Kuwait: options and challenges. *Clean. Technol. Environ. Policy* 14 (4), 619–624. doi:10.1007/s10098-011-0421-x

Erdmann, J. (2015). *Sumo's emission models*. Berlin/Heidelberg, German: Springer. Available online at: https://sumo.dlr.de/docs/Models/Emissions.html.

Gabler, V., and Wollherr, D. (2024). Decentralized multi-agent reinforcement learning based on best-response policies. Front. Robotics AI 11, 1229026. doi:10.3389/frobt.2024.1229026

Gao, J., Shen, Y., Liu, J., Ito, M., and Shiratori, N. (2017). Adaptive traffic signal control: deep reinforcement learning algorithm with experience replay and target network. *arXiv Prepr. arXiv:1705.02755*. doi:10.48550/arXiv.1705.02755

Gartner, N., Messer, C. J., and Rathi, A. (2001). Traffic flow theory: a state-of-the-art report. *Transp. Res. Board Special Rep.* 165. Available online at: https://www.researchgate.net/publication/248146380_Traffic_Flow_Theory_A_State-of-the-Art_R eport.

 $Helbing, D., Farkas, I., and Vicsek, T. (2000). Simulating dynamical features of escape panic. \\ \textit{Nature 407 (6803)}, 487-490. \\ \textit{doi:} 10.1038/35035023$

IQAir (2023). Air quality in Kuwait city. Available online at: https://www.iqair.com/us/kuwait.

Jarrah, A. A., Haymoor, Z. S., Al-Masri, H. M. K., and Almomany, A. M. (2022a). High-performance implementation of power components on fpga platform. *J. Electr. Eng. Technol.* 17 (3), 1555–1571. doi:10.1007/s42835-022-01005-6

Jarrah, A. A., Almomany, A. M., and Haymoor, Z. S. (2022b). High-performance implementation of wideband coherent signal-subspace (css) based doa algorithm on fpga. *J. Electr. Eng. Technol.* 17 (6), 2831–2846. doi:10.1007/s42835-022-01098-z

Jiang, X., Ma, X., Adnan, M., and Zhang, Y. (2017). Evaluating effects of traffic signal optimization on vehicle emissions and fuel consumption. *Transp. Res. Part D* 54, 282-290. doi:10.1080/03081060.2011.651877

Kouvelas, A., Aboudolas, K., Papageorgiou, M., and Diakaki, C. (2017). A hybrid strategy for real-time traffic signal control of urban road networks. *Transp. Res. Part C* 78, 84–96. doi:10.1109/TITS.2011.2116156

Li, L., Lv, Y., and Wang, F. (2014). Traffic signal timing via deep reinforcement learning. $IEEE/CAA\ J.\ Automatica\ Sinica\ 3$ (3), 247–254. doi:10.1109/jas.2016.7508798

Lin, S., De Schutter, B., Xi, Y., and Hellendoorn, H. (2015). Efficient network-wide model-based predictive control for urban traffic networks. *Transp. Res. Part C* 57, 127-147. doi:10.1016/j.trc.2012.02.003

List of countries and territories List of countries and territories by motor vehicles per capita, Wikipedia, Global, India, France motorization rates (2024). Available online at: https://en.wikipedia.org/wiki/List_of_countries_and_territories_by_motor_vehicles_per_capita.

Maschi, F., Alonso, G., Hock-Koon, A., et al. (2021). From research to proof-of-concept: analysis of a deployment of fpgas on a commercial search engine. arXiv Prepr. Tech. Rep. doi:10.3929/ethz-b-000501533

Michailidis, P., Michailidis, I., Lazaridis, C. R., and Kosmatopoulos, E. (2025). Traffic signal control via reinforcement learning: a review on applications and innovations. *Infrastructures* 10 (5), 114. doi:10.3390/infrastructures10050114

Muralidharan, A., Pedarsani, R., and Varaiya, P. (2015). Analysis of fixed-time control. Transp. Res. Part B Methodol. 73, 81–90. doi:10.1016/j.trb.2014.12.002

Othman, K., Wang, X., Shalaby, A., and Abdulhai, B. (2025). Multimodal adaptive traffic signal control: a decentralized multi-agent reinforcement learning approach. *Multimodal Transp.* 4 (1), 100190–190. doi:10.1016/j.multra.2025.100190

Papageorgiou, M., Diakaki, C., Dinopoulou, V., Kotsialos, A., and Wang, Y. (2003). Review of road traffic control strategies. *Proc. IEEE* 91 (12), 2043–2067. doi:10.1109/jproc.2003.819610

Pham, H. D., Narasimhamurthy, S. M., Mehran, M. B., Manley, E., and Ahmed, A. (2025). Reinforcement learning based estimation of shortest paths in dynamically changing transportation networks. *Front. Future Transp.* 6, 1524232. doi:10.3389/ffutr.2025.1524232

Qasaimeh, M., Denolf, K., Lo, J., Vissers, K., Zambreno, J., and Jones, P. H. (2019). "Comparing energy efficiency of cpu, gpu and fpga implementations for vision kernels," in 2019 IEEE International Conference on embedded software and systems (ICESS), 1–8. doi:10.1109/ICESS.2019.8782524

Qin, Y., Luo, Q., and Xiao, T. (2024). Capacity modeling for mixed traffic with connected automated vehicles on minor roads at priority intersections. *Transp. Plan. Technol.* 0 (0), 1–25. doi:10.1080/03081060.2024.2428410

Ramamoorthy, R. (2025). Applicability and assessment of diverse fpga architectures and algorithms for the internet of things applications. Available online at: https://www.academia.edu/79354820/Field_Programmable_Gate_Array_FPGA_Based_IoT_for_Smart_City_Applications.

Saadi, A., Abghour, N., Chiba, Z., Moussaid, K., and Ali, S. (2025). A survey of reinforcement and deep reinforcement learning for coordination in intelligent traffic light control. *J. Big Data* 12, 84. doi:10.1186/s40537-025-01104-x

Tan, K. L., Tao, F., Zhao, Z., and Chang, Y. (2024). "Deep reinforcement learning for adaptive traffic signal control," doi:10.48550/arXiv.1911.06294

Van der Pol, E., and Oliehoek, F. A. (2016). "Coordinated deep reinforcement learners for traffic light control," in NIPS Workshop on learning, Inference and control of multi-agent systems.

Varaiya, P. (2013). Max pressure control of a network of signalized intersections. Transp. Res. Part C 36, 177–195. doi:10.1016/j.trc.2013.08.014

Wang, X., Taitler, A., Sanner, S., and Abdulhai, B. (2024). "Mitigating partial observability in adaptive traffic signal control with transformers," in *Proceedings of the TRC-30 Conference on emerging technologies in transportation systems*. Available online at: https://arxiv.org/abs/2409.10693v1.

Wongpiromsarn, T., Uthaicharoenpong, T., Wang, Y., Yip, N., and Hsieh, H. (2012). "Distributed traffic signal control for maximum network throughput," in 13th Int. IEEE Conf. On intelligent transportation systems.

Wu, J., Ghosal, D., Zhang, M., and Chuah, C.-N. (2018). Delay-based traffic signal control for throughput optimality and fairness at an isolated intersection. *IEEE Trans. Veh. Technol.* 67 (2), 896–909. doi:10.1109/TVT.2017.2760820

Xiao, F. (2025). Advances in reinforcement learning for traffic signal control. Intell. transp. infrastruct. 4. doi:10.1093/iti/liaf009