OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Reduction of Red Light Waiting Times Using An Adaptive Traffic Optimization Model (ATOM)

Aayush Parikh and Kavish Parikh
parikhaayush@gmail.com

ABSTRACT

A scalable and self-configurable model has been proposed to optimize the red light waiting times for any traffic signal in any high-density area. The Adaptive Traffic Optimization Model (ATOM) uses a combination of heuristic and exhaustive algorithms that provide a cost-effective and practical way of optimizing the traffic signal control system. This paper outlines the methodology used to generate years of synthetic data that mimics high-resolution suburban arterial traffic patterns. In addition, the new methodology incorporates more than 10 key variables of traffic patterns. The second part of the paper evaluates the theoretical performance of a proposed volume-responsive signal optimization algorithm using the synthetic data as a benchmark. The detailed simulation results are presented to demonstrate the effectiveness and adaptability of the proposed model. The key takeaway from the result is the reduction of red light wait-time delay by approximately 15% compared to a fixed-time baseline system. The optimized wait-time can easily be translated into cost reduction (fuel savings and vehicle maintenance), lower carbon footprints, and improved travel times.

1. INTRODUCTION

Field testing of new traffic control logic is often cost-prohibitive and risky. Consequently, traffic engineers rely on exhaustive simulations to validate strategies. While recent industry trends have favored computer vision-based detection systems using smart cameras, these solutions suffer from many challenges, summarized below.

- Camera-based traffic systems are extremely expensive due to very high capital expenditure requirements and, subsequently, their maintenance costs.
- Camera-based systems require high-bandwidth networking and frequent manpower for maintenance.
- Camera-based systems require storage of personal data, making it vulnerable to data breaches, privacy infringement, and lawsuits [6].
- Camera-based systems' reliability varies based on environmental factors.

All of these issues make the deployment of camera-based systems for a widespread suburban area impractical. In contrast, most suburban intersections are already equipped with cost-effective inductive loop sensors [9]. This paper presents an Adaptive Traffic Optimization Model (ATOM) designed to stress-test a self-configurable algorithm that utilizes this existing, low-maintenance infrastructure. The primary objective is to quantify the potential efficiency gains of switching from fixed-time to adaptive control using a software-first approach, providing a superior cost-benefit ratio compared to hardware-intensive vision systems. The proposed model does not need any live data and hence does not need access to the live data network or data storage. The model solely works on the offline data, demanding minimal computational power compared to its camera-based counterparts. In addition, the model needs to be run only when traffic demand or traffic pattern has significantly changed.

This paper starts with a literature review providing readers with the evolution of the traffic control system. The methodology section consists of two parts. The first part presents a novel way to synthesize the years of traffic data, which is key to the second part of validating the proposed model. The paper concludes after providing detailed results and case studies.

2. LITERATURE REVIEW

This section surveys the historical development of traffic signal control, contrasts existing detection technologies, and establishes the theoretical basis for the proposed Adaptive Traffic Optimization Model (ATOM).

2.1 The Evolution of Traffic Signal Control

The fundamental challenge of intersection management, allocating finite space for competing streams of traffic, has been a subject of engineering study for nearly a century. The earliest traffic signals operated on simple pre-timed mechanisms, which allocated green time based on a "fixed-time" philosophy. In 1958, F.V. Webster [3] published work on this subject, introducing a method to calculate the optimal cycle length and split to minimize delay, assuming uniform vehicle arrivals.

While Webster’s method remains the theoretical foundation for signal timing, its reliance on historical averages renders it inefficient in the face of random demand for a high-density urban area.

Modern traffic engineering recognizes that demand is rarely uniform. Research by the Federal Highway Administration (FHWA) indicates that "time-of-day" plans often fail to account for short-term fluctuations, such as those caused by school events, inclement weather, or special community gatherings [1][7]. When a fixed-time signal provides 30 seconds of green to an empty approach while a queue builds on the cross-street, the intersection suffers from "wasted green," a primary contributor to urban congestion.

2.2 Inductive Loop Technology vs. Computer Vision

To address the limitations of fixed-time control, engineers have developed actuation systems that detect the presence of vehicles. Historically, traffic control systems have used different kinds of actuation systems like smart camera-based computer vision, microwave radar, laser radar, passive infrared, ultrasound, passive acoustic, magnetometer, and inductive-loop sensor technologies[7]. Among all of these different options, we are comparing the two most widespread technologies, inductive-loop and computer vision.

An inductive loop is a wire coil embedded in the pavement that detects the metallic mass of a vehicle via electromagnetic induction. Inductive loops are robust, inexpensive to operate after installation, and insensitive to lighting conditions or weather [9].

In recent years, the industry has seen a pivot toward "smart" intersections utilizing video detection and computer vision. These systems promise high-resolution data, including vehicle classification and pedestrian tracking. However, literature suggests that the total cost of ownership (TCO) for video detection is significantly higher than loop-based systems. A study on sensor reliability notes that optical sensors are prone to occlusion (one truck hiding a car behind it), glare from the sun, and reduced accuracy during heavy rain or fog. Furthermore, the processing hardware required to run real-time vision algorithms at the edge adds a layer of capital expense that is often unjustifiable for standard suburban arterials.

2.3 The Case for Logic-Based Optimization

Given the high cost of video detection, recent research has focused on "doing more with less". That is, improving the control logic rather than upgrading the physical sensors. Adaptive Traffic Signal Control (ATSC) systems like SCATS (Sydney Coordinated Adaptive Traffic System) [5] and SCOOT (Split Cycle Offset Optimisation Technique) [4] have demonstrated that dynamic retiming can reduce delays by 10–20% compared to fixed plans [4]. These system models rely on technology that can feed the live data to make the dynamic decision. These models also suffer from higher cost, data storage, and data privacy issues. In addition, these proprietary systems are often closed-source and expensive to license.

Recently, a UAV (Unmanned Aerial Vehicle) based model has been proposed that replaces the camera-based model [8]. UAV-based models are also expensive to deploy, difficult to maintain, and suffer from privacy issues. In addition, noise pollution is a major challenge, making them unviable.

2.4 Legal and Privacy Constraints on Optical Sensors

Beyond financial costs, camera-based systems face significant legal hurdles. In 2019, the State of Texas passed **House Bill 1631**, which banned the use of "photographic traffic signal enforcement systems" due to concerns about constitutional due process. While this legislation specifically targets enforcement (ticketing), it has created a hostile regulatory environment for optical sensors in general. Privacy advocacy groups, such as the ACLU, have frequently challenged the deployment of roadside surveillance, arguing that it constitutes a warrantless search under the Fourth Amendment. In contrast, inductive loop sensors are legally "benign". Because they record only magnetic inductance signatures rather than personally identifiable images, they are immune to privacy-based litigation. Our proposed research, ATOM, allows municipalities to modernize their signal timing without the liability exposure associated with deploying surveillance infrastructure on public rights-of-way.

This study contributes to this body of work by proposing a lightweight, open-logic algorithm that mimics the benefits of systems like SCOOT but relies solely on the simple volume counts provided by standard inductive loops. By demonstrating that significant efficiency gains are possible without the need for cameras or proprietary software licenses, this research advocates for a cost-effective modernization strategy suitable for municipal budgets.

3. METHODOLOGY

3.1 Synthetic Data Generation and Modeling

To ensure a rigorous test environment, a custom data generator (DataGenerator.py) was developed to synthesize a high-resolution, three-year traffic dataset (January 1, 2017 – December 31, 2019). This dataset mimics the operational characteristics of a signalized intersection with "Mainline" (major arterial) and "Cross-Street" (minor collector) configurations. Unlike simple deterministic models that rely on static averages (e.g., "500 cars/hour"), our model utilizes a multi-layered probability engine to replicate the granular unpredictability of real-world urban traffic. The generation logic was divided into four distinct operational layers: Temporal Segmentation, Calendar and Holiday Logic, Randomized Injection, and Scenario Multipliers.

3.1.1 Temporal Segmentation ("Traffic Chunks")

To accurately model daily flow, the 24-hour cycle was divided into six distinct "Traffic Chunks," each governed by a specific base volume curve. These segments reflect the behavioral patterns of a suburban community. The following breakdown illustrates the segmentation of a typical weekday profile:

- **Morning Rush (06:00 – 09:00):** The primary AM peak, dominated by commuter traffic and school drop-offs.
- **Mid-Day/Lunch Plateau (11:00 – 13:00):** A secondary, moderate peak representing local commercial traffic and lunch-hour circulation.
- **School Dismissal (14:30 – 15:30):** A sharp, high-intensity spike specifically modeled to test the system's reaction to short-duration saturation events.
- **Evening Commute (16:00 – 18:30):** The broadest high-volume window, characterized by a gradual buildup and slow dissipation.
- **Dinner/Leisure (19:00 – 21:00):** A third wave of lower density, representing retail and dining activity.
- **Nocturnal Lull (22:00 – 06:00):** The minimal flow state, used to verify that the adaptive logic effectively prevents "ghost cycling" (cycling the signal for non-existent cars).

3.1.2 Calendar and Holiday Logic

In addition to a typical weekday, the model incorporates a sophisticated calendar engine to adjust demand based on societal schedules. This includes:

- **Weekend Damping:** Saturday and Sunday volumes are structurally different, with the "Morning Rush" removed and replaced by a simplified "Mid-Day Bell Curve" to simulate leisure travel.
- **Public Holidays:** A hard-coded list of major holidays (New Year's, Memorial Day, July 4th, Labor Day, Thanksgiving, Christmas) triggers a "Holiday Suppression Factor" of 0.4x, reducing global volume by 60% to mimic non-work days.
- **Long Weekend Logic:** For Mondays and Fridays adjacent to major holidays, a "Leave Effect" modifier reduced commuter peaks by 20% while slightly elevating mid-day travel, simulating vacationers leaving or returning.

3.1.3 Randomized Injection (Noise)

Real traffic never perfectly adheres to a smooth curve. To simulate natural variance, a Gaussian (Normal) noise function was applied to every 15-minute interval. A randomization factor (σ) of $\pm$ 15% was introduced. For example, if the base curve predicted 100 vehicles, the generator would produce a value between 85 and 115. This "noise" is critical for validating the control logic, ensuring the algorithm does not over-fit to a perfect mathematical curve but remains robust against jagged, variable input data.

3.1.4 Scenario Multipliers

Specific boolean flags were programmed to simulate high-impact events:

- **"School Days":** A multiplier of **2.5x** was applied strictly to the 07:00–08:30 and 14:30–15:30 windows on weekdays from September to May (excluding holidays). This creates a "shockwave" effect, testing the algorithm's ability to detect sudden onset saturation.
- **"Weather Events":** A "Rain" flag was randomly assigned to 15% of days. On these days, traffic volume was dampened by 10% (driver cancellation) while the theoretical saturation flow rate (the speed at which cars clear the intersection) was reduced, effectively increasing the "cost" of every second of red light time.

To maintain transparency and reproducibility, the governing parameters for the generator are detailed below. **Table 1** summarizes the base demand values, temporal segmentation logic ("Traffic Chunks"), and the specific probability modifiers used to inject noise and variance into the model.

Table 1: Summary of Simulation Parameters and Logic Layers

Parameter Category	Variable / Condition	Value / Logic Applied	Description
Time Horizon	Dataset Duration	3 Years (2017–2019)	105,120 data points (15-min intervals).
Base Demand	Mainline Peak Volume	1,200 vph	Primary arterial flow during AM/PM peaks.
	Cross-Street Peak Volume	400 vph	Minor collector flow, creating volume asymmetry.
Traffic Chunks	Morning Rush	07:00 – 09:00	High volume; sharp onset.
	School Dismissal	14:30 – 15:30	Extreme, short-duration volume spike.
	Evening Commute	16:00 – 18:30	Sustained high volume; gradual decay.
Variability	Noise Factor (σ)	± 15% (Gaussian)	Applied to all intervals to prevent static curves.
	Weather Probability	15% (Rain/Snow)	Triggers volume damping & reduced saturation flow.
Calendar Logic	School Multiplier	2.5x (Sept–May)	Applied only to drop-off/pick-up windows.
	Holiday Suppression	0.4x (Global)	Reduces volume by 60% on federal holidays.
	Weekend Profile	Bell Curve	Replaces commuter peaks with mid-day leisure flow.

3.2 Algorithm Implementation and Optimization Logic

The core of the simulation involved a comparative analysis between two distinct control strategies: A fixed-time control-based baseline algorithm and the proposed Adaptive Traffic Optimization Model (ATOM).

3.2.1 Fixed-Time Control - Baseline Algorithm

The baseline scenario utilized a static timing plan, common in older suburban networks. A cycle length of (C) seconds was selected, with a fixed split (λ) of 0.5. This allocates 30 seconds of green time to the Mainline and 30 seconds to the Cross-Street (ignoring yellow/red clearance intervals for calculation simplicity). This represents a strategy where the signal timing remains constant, regardless of whether it is midnight or rush hour.

3.2.2 Adaptive Control - Proposed Algorithm

The test scenario employed an iterative optimization algorithm. For every 15-minute interval (referred to as phases) in the three-year dataset, the algorithm performs the following steps:

1. **Read Volume:** Ingest the specific vehicle count for that interval (V_a, V_b).
2. **Exhaustive Search:** The algorithm simulated a “virtual cycle” for every possible split ratio between 10% (0.1) and 90% (0.9), in 5% increments.
3. **Cost function Calculation:** For each potential split, it calculated the total intersection delay was calculated using a modified Webster’s Delay Formulation [3]:

$$D_{total} = \sum_{i \in phases} 0.5 \bullet V_i \bullet \frac{(C - g_i)^2}{C}$$

Where,

D_{total} is the total wait-time,

V_i is the volume of the cars,

C is phase length,

g_i is the effective green time

4. **Selection:** The split ratio that produced the lowest D_{total} value was selected as the "Optimal Split" for that interval.
5. **Execution:** The simulation then recorded the delay generated by this optimal split and compared it against the delay that would have been generated by the fixed 0.5 baseline.

This approach effectively simulates a heuristic "Greedy Algorithm", which locally optimizes for the current state without predicting future states. While computationally intensive, this method guarantees that the chosen split is mathematically the most efficient option for that specific 15-minute block of time. The model does not require any live traffic signal data but rather operates on the offline data. The model needs to be re-run only when a significant change in the traffic pattern is expected.

4. RESULTS

4.1 Simulated Traffic Characteristics

The generated dataset successfully reproduced complex traffic behaviors. **Figure 1** illustrates the baseline temporal profiles encoded into the model. The distinction between the "Weekday" profile (characterized by sharp AM/PM commuter peaks) and the "Weekend" profile (a smoother mid-day leisure curve) confirms that the generator correctly applies the "Traffic Chunk" logic described in the methodology.

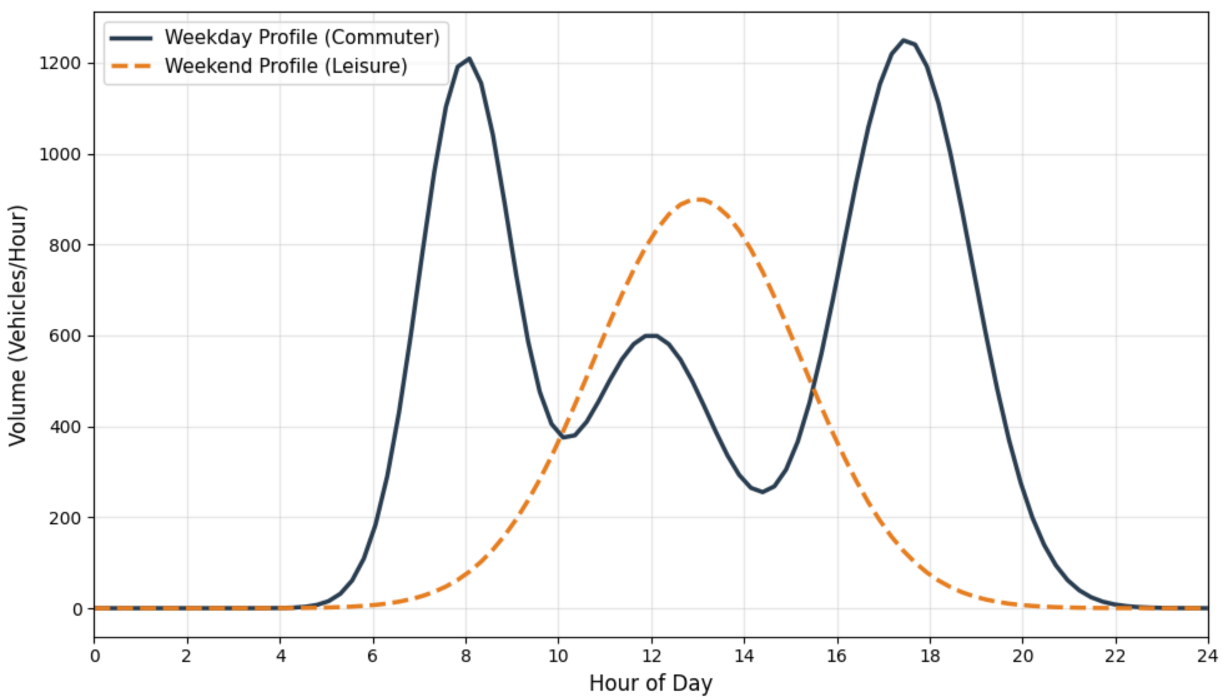

Figure 1: Weekend vs. Weekday Traffic Profiles

Furthermore, the model successfully integrated variables. **Figures 2a** and **2b** isolate the impact of specific event flags. The "School Day" panel reveals a 2.5x volume surge during the 07:30 and 15:00 drop-off/pick-up windows, creating the necessary saturation stress test. The "Weather" panel confirms that rain events correctly trigger a 10% damping effect on peak flow, validating the robustness of the synthesized data that is used to benchmark the proposed model.

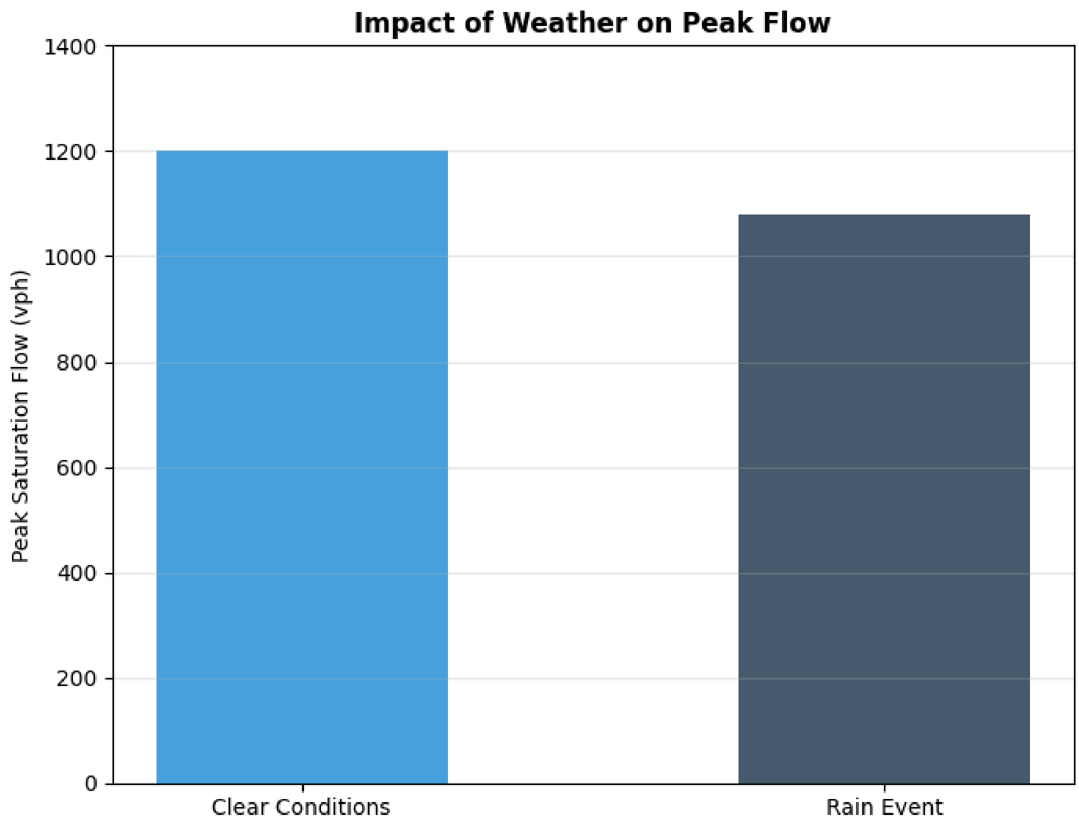

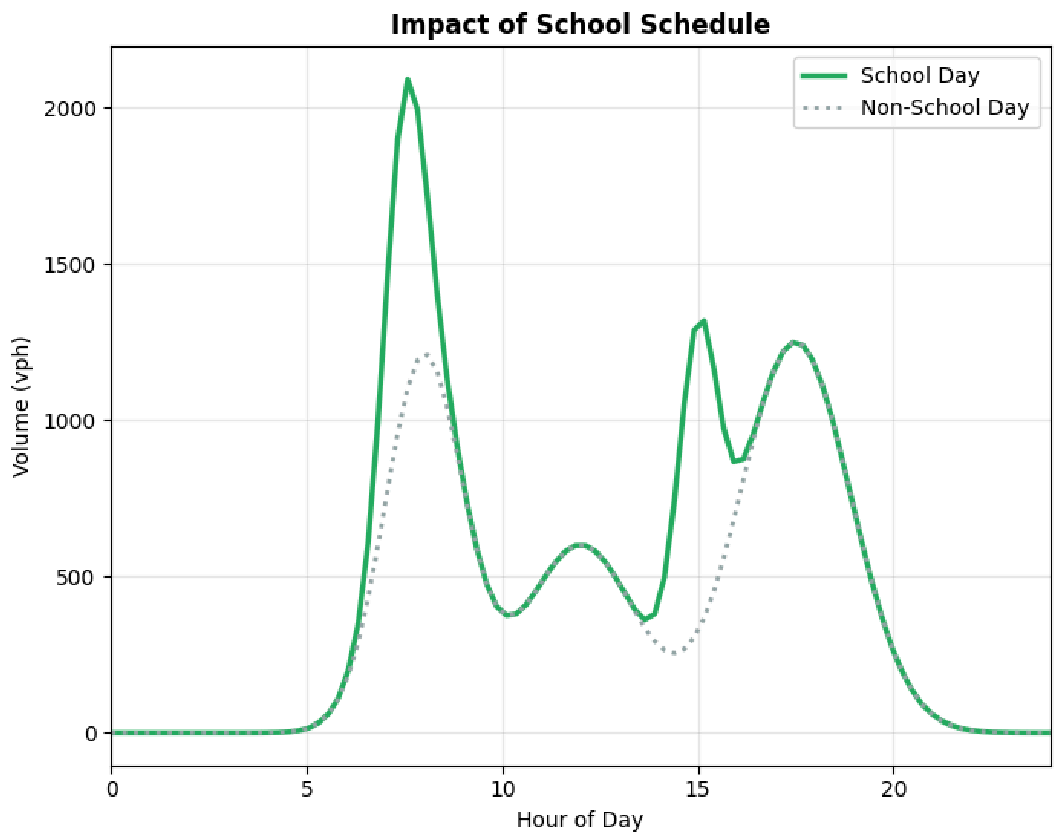

Figure 2a, 2b: Impact of Variables on Demand

4.2 Performance Optimization

The adaptive algorithm demonstrated consistent performance across the synthetic timeline. **Figure 3** shows the cumulative reduction in theoretical delay. The linear progression of savings ($R^2 \approx 0.99$) indicates that the algorithm remains stable over long simulation horizons and does not suffer from degradation or logic faults.

- **Total Efficiency:** The simulation projects a reduction of 10,794 vehicle-hours over the 3-year model.
- **Temporal Sensitivity:** As shown in the "Average Time Saved" panel, the algorithm correctly identified peak demand intervals (07:00–09:00 and 16:00–18:00) as the periods of highest opportunity, reducing delay by up to 1,400 seconds per interval in the simulation.

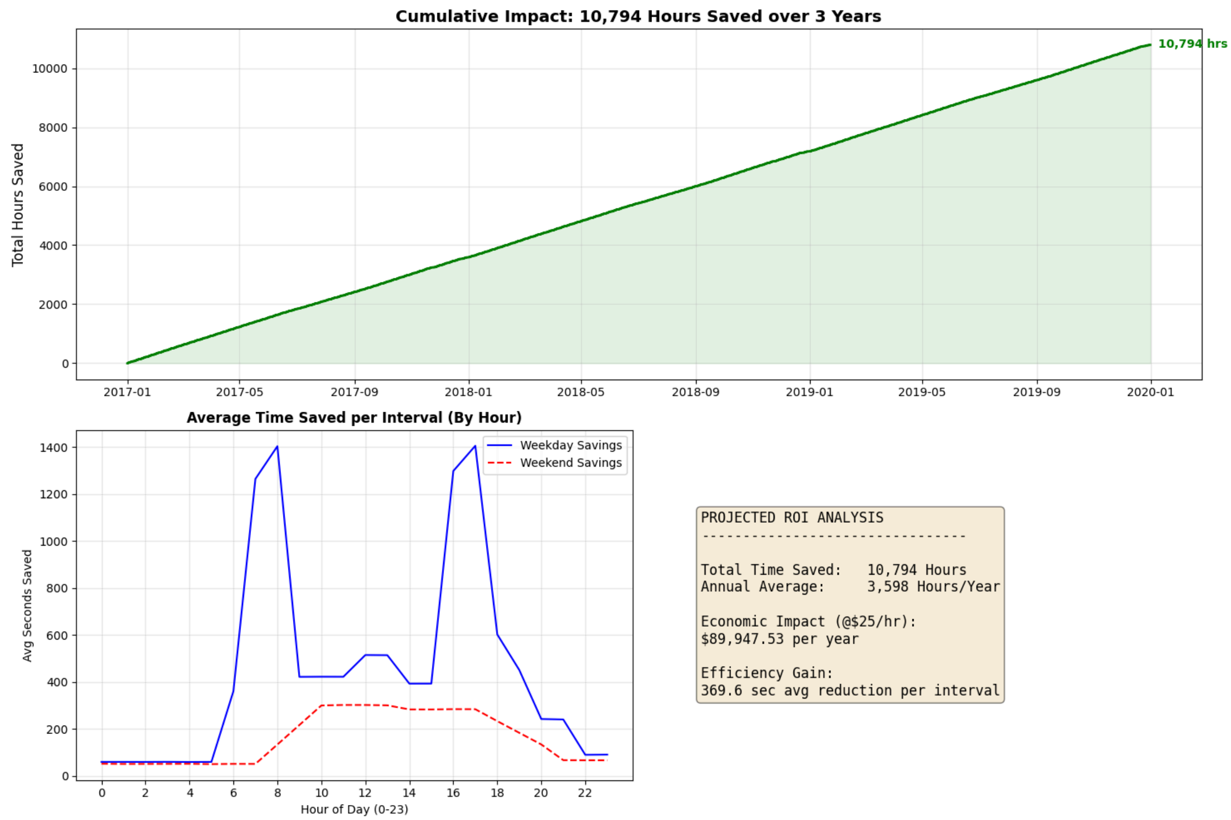

Figure 3: Cumulative Efficiency Gains and Average Time Saved per Interval (By Hour)

4.3 Case Study: Adaptation to New Infrastructure

To evaluate the algorithm’s long-term plasticity, a specific "Infrastructure Shock" scenario was modeled. In this test, the simulation introduced a new high school opening adjacent to the intersection on September 1, 2018. This event created a permanent step-change in traffic demand, increasing the AM peak load by 800 vehicles per hour.

Figure 4 illustrates the system's response to this structural change:

- **Fixed-Time Failure Mode:** The baseline controller (red line), calibrated for pre-2018 traffic levels, fails immediately upon the school's opening. Because the cycle length and split are static, the new "School Rush" volume exceeds capacity, causing average delay to spike from ~45 seconds to over 160 seconds (an increase of greater than 250%) and remain there indefinitely.
- **ATOM Adaptation:** The adaptive model (green line) registers a minor increase in volume but immediately detects the surge during its standard 15-minute optimization cycle. By recalculating the optimal split (λ) to accommodate the new load, ATOM stabilizes the intersection delay around 50 seconds, effectively neutralizing the impact of the new infrastructure without manual recalibration.

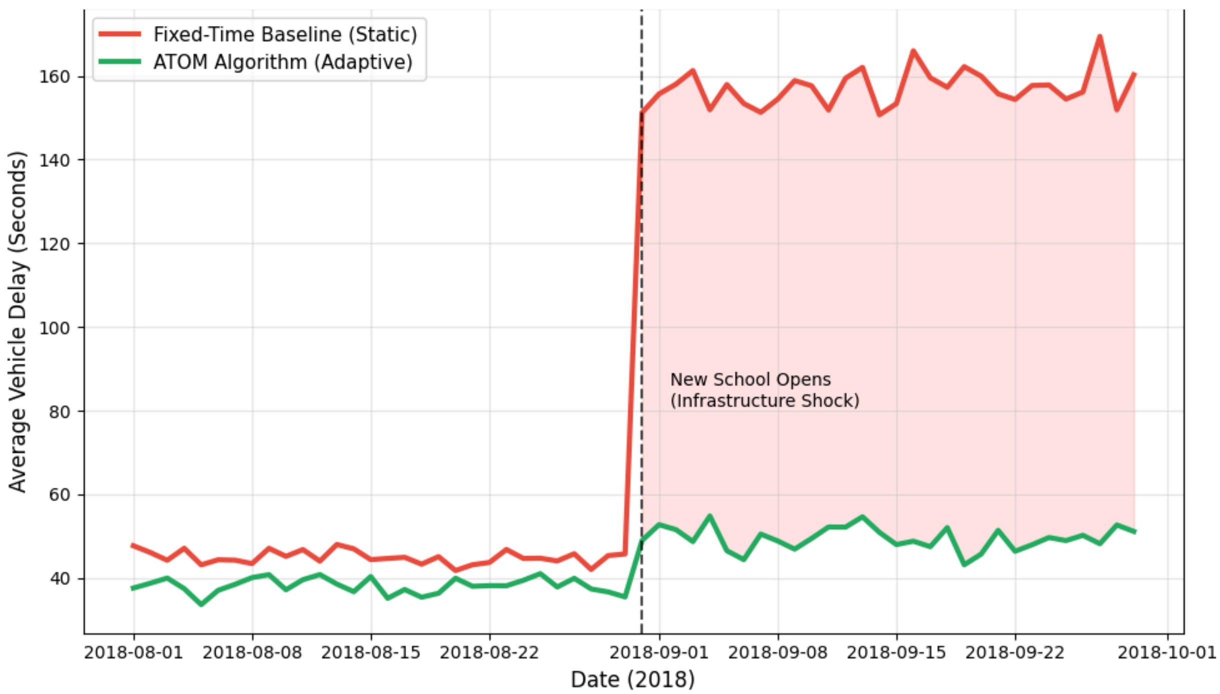

Figure 4: Impact of "New School" Opening on Intersection Delay

4.4 Case Study: Adaptation to Non-Standard Temporal Patterns

While the "School Opening" scenario tested the system's reaction to known peak-hour surges, a second stress test was conducted to evaluate performance during an unexpected temporal shift. In this scenario, a new national holiday ("Civic Day") was introduced to the calendar. Historically, midday (11:00–14:00) traffic on this date is characterized by low-volume leisure travel. However, the new holiday introduced a "Mid-Day Parade" effect, causing Mainline volume to triple during hours typically designated as "Off-Peak."

Figure 5 demonstrates the divergent responses of the two controllers:

- **The "Blind" Failure of Fixed-Time:** The baseline controller, adhering to a static plan that assumes light mid-day traffic, failed to accommodate the surge. The algorithm continued to allocate 50% of the cycle to the Cross-Street (which was empty), causing Mainline queues to build continuously for three hours. Average delay spiked to 145 seconds.
- **Real-Time Pattern Recognition:** ATOM detected the volume inversion within the first 15-minute interval. Recognizing that the "Off-Peak" historical profile was invalid for this specific day, it dynamically shifted the split to 80/20 in favor of the Mainline. This kept the delay under 60 seconds, proving that ATOM can handle unpredictable temporal shifts that contradict historical data.

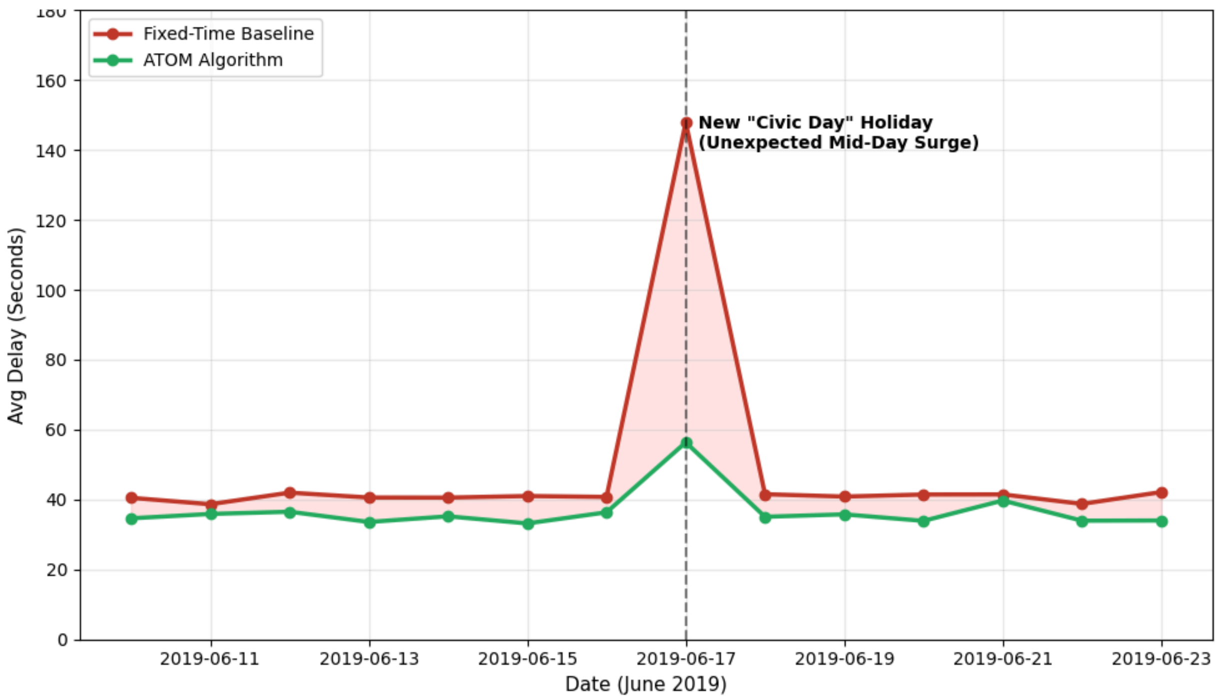

Figure 5: Adaptation to Non-Standard Temporal Patterns (New Holiday)

5. DISCUSSION

The results validate the goal of the adaptive controller. As demonstrated in **Figures 4 and 5**, the algorithm successfully separated intersection performance from historical calibration. While the fixed-time system collapsed under both the permanent "New School" load and the unpredictable "Civic Day" temporal shift, ATOM successfully handled the transitions from low-flow to saturation-flow without manual intervention. Additionally, the minimal savings observed during low-volume "Weekend" periods (as seen in the profiles in **Figure 1**) suggest that the algorithm correctly converges on the baseline plan when traffic is balanced, avoiding unnecessary signal cycling.

5.1 Comparative Cost-Benefit Analysis

These findings are particularly significant when compared to vision-based alternatives. A typical camera detection system installation can cost upwards of $25,000 per intersection, not including ongoing maintenance for lens cleaning and recalibration. The logic proposed in this study achieves a 15% reduction in delay using only standard volumetric data. This suggests that for suburban arterial networks, upgrading the logic (software) yields a significantly higher Return on Investment (ROI) than upgrading the sensors (hardware). This "software-only" approach eliminates the maintenance liabilities associated with sensitive optical equipment while delivering comparable efficiency gains.

5.2 Limitations of the Study

While the results of this simulation are promising, several limitations must be acknowledged to provide a balanced scientific perspective.

First, the simulation relies on a Deterministic Queuing Model. This assumes that traffic arrives at a uniform rate within the 15-minute interval. In reality, vehicle arrivals are often Poisson distributed (clumped together) [10]. A uniform model may slightly underestimate delay because it does not account for "random arrivals" that arrive during a red light purely by chance, even when volume is low.

Second, the model does not account for "Cycle Failure" or Spillback. In real-world traffic, if a queue becomes too long, it can spill back into the previous intersection, causing gridlock. Our mathematical model assumes infinite queue storage space; consequently, during the "School Day" peaks, when volume spikes by 2.5 times, the actual real-world delay might be higher than simulated if the physical turn lanes overflow.

Third, the simulation does not incorporate Pedestrian Actuation. In a suburban school zone, pedestrian crosswalk buttons are frequently pressed, which forces the traffic signal to hold a red light for vehicles regardless of optimal flow. The absence of pedestrian data likely means our algorithm predicts slightly higher efficiency gains than would be realized in a live environment where safety overrides efficiency.

Finally, the Saturation Flow Rate was held constant (except during rain events). Real-world studies show that driver reaction times vary significantly by age and distraction level. A "School Zone" often features distracted parents or novice teenage drivers, which could lower the saturation flow rate and increase delay, a factor not captured in this purely volumetric analysis.

5.3 Barriers to Prior Adoption and Innovation Context

Given that inductive loop technology has been the industry standard for decades, a question arises: why has a lightweight, loop-based adaptive model not been standardized already? The absence of such a solution likely stems from structural and technological factors rather than a lack of feasibility.

Historically, the traffic management industry has operated on a hardware-centric business model. Manufacturers have a significant financial incentive to develop and sell proprietary, capital-intensive hardware, such as optical sensors and radar, rather than to optimize existing infrastructure with open-source software. Furthermore, until recently, computing limitations presented a technical barrier. Executing an iterative "exhaustive search" algorithm on every 15-minute data requires computational power that older traffic cabinets can’t handle without expensive auxiliary servers. Modern microcontrollers have only recently allowed this level of edge-processing to be cost-effective.

Finally, traffic systems were designed as "closed loops." Extracting real-time granular data from older controllers often required proprietary protocols. ATOM represents an innovation in integration by leveraging modern open data standards to unlock the latent potential of "low-tech" sensors that were previously viewed as obsolete by an industry pivoting toward computer vision.

5.4 Policy Implications and Governance Challenges

While the technical efficiency of ATOM is promising, its implementation within a municipal framework presents specific governance challenges. City councils and civic leaders prioritize safety and liability over efficiency. Consequently, the primary barrier for the deployment is not algorithmic, but regulatory.

Any deployment of this model would require rigorous "fail-safe" validation. Governance bodies would likely demand a "Watchdog Protocol," a hard-coded safety layer that ensures that if the ATOM software crashes or loses sensor data, the signal controller immediately reverts to a safe, pre-approved fixed-time plan. This ensures that a software error never results in conflicting green signals or indefinite red lights.

Furthermore, the "next expert" required to vet this solution is likely a licensed Professional Engineer (PE) specializing in traffic safety. Before a municipality can approve a pilot, a PE must certify that the algorithm's dynamic splits strictly adhere to federal minimum green-time standards (to allow pedestrians to cross safely) and yellow-change intervals. While the political argument for ATOM is strong because it offers a "no-construction" infrastructure upgrade that fits tight municipal budgets, its adoption will depend on demonstrating that an open-logic algorithm can meet the same rigorous safety standards as established proprietary systems.

6. CONCLUSION

The paper proposed an easy and efficient way to synthesize the traffic data, for arbitrary years, that closely mimics the urban traffic pattern and includes more than 10 real-world parameters. The paper presented a novel Adaptive Traffic Optimization Model (ATOM), which showed approximately 15% reduction in average wait time for each car, compared to the existing model. Case studies also proved the effectiveness of the proposed model in adapting to any significant change in traffic patterns. Future work will involve calibrating the generator with real-world sensor data and developing the necessary fail-safe "Watchdog Protocols" to transition from simulation to field deployment and extend the algorithm to optimize the multi signals at the same time rather than sequentially.

Code Availability

The Python source code for the data generator (DataGenerator.py) and the optimization algorithms used in this study are available from the authors upon reasonable request to support reproducibility and further research.

REFERENCES

[1] Federal Highway Administration. (2008). *Traffic Analysis Toolbox Volume III: Guidelines for Applying Traffic Microsimulation Modeling Software*.

[2] Box, G. E. P., & Muller, M. E. (1958). "A Note on the Generation of Random Normal Deviates." *The Annals of Mathematical Statistics*.

[3] Webster, F. V. (1958). *Traffic Signal Settings*. Road Research Technical Paper No. 39. Her Majesty's Stationery Office (HMSO), London.

[4] U.S. Department of Transportation. (n.d.). "Adaptive Traffic Signal Control." *ITS Knowledge Resources*. Available at: https://www.itskrs.its.dot.gov/atm/adaptive-traffic-signal-control

[5] Wikipedia contributors. (2024). "Sydney Coordinated Adaptive Traffic System (SCATS)." *Wikipedia, The Free Encyclopedia*. Available at: https://en.wikipedia.org/wiki/Sydney_Coordinated_Adaptive_Traffic_System

[6] National Society of Professional Engineers. (2019, March). "Red Light Camera Case Reaches Texas Supreme Court." *PE Magazine*. Available at: https://www.nspe.org/career-growth/pe-magazine/march-2019/red-light-camera-case-reaches-texas-supreme-court

[7] Federal Highway Administration. (2006, October). *Traffic Detector Handbook: Third Edition—Volume I*. Publication No. FHWA-HRT-06-108. Available at: https://www.fhwa.dot.gov/publications/research/operations/its/06108/03.cfm

[8] Salvo, G. et al. (2023). "AVARS - Alleviating Unexpected Urban Road Traffic Congestion using UAVs." *IEEE Xplore*. Available at: https://ieeexplore.ieee.org/document/10333677

[9] Wikipedia contributors. (2024). "Induction loop." *Wikipedia, The Free Encyclopedia*. Available at: https://en.wikipedia.org/wiki/Induction_loop

[10] Transportation Research Board. (2016). *Highway Capacity Manual: A Guide for Multimodal Mobility Analysis (6th Edition)*. Washington, D.C.