CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

CCAT-CTS 26-06
March 2026

Cybersecurity of next-generation traffic signal control with CAVs

Raphael Stern
Amirhossein Kiani

Center for Transportation Studies

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

Report No. CTS 26-06

March 2026

Project Start Date: September 1, 2023

Project End Date: December 31, 2025

Cybersecurity of next-generation traffic signal control with CAVs

by
Raphael Stern
Associate Professor
University of Minnesota

Amirhossein Kiani
Research Assistant
University of Minnesota

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

DISCLAIMER

Funding for this research was provided by the Center for Connected and Automated Transportation under Grant No. 69A3552348305 of the U.S. Department of Transportation, Office of the Assistant Secretary for Research and Technology (OST-R), University Transportation Centers Program. The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the Department of Transportation, University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Suggested APA Format Citation: Stern, R. and Kiani, A. (2026). *Cybersecurity of next-generation traffic signal control with CAVs*, CTS 26-06.

Contacts

For more information:

Raphael Stern
500 Pillsbury Dr SE
Phone: (612) 625-6354
Email: rstern@umn.edu

CCAT
University of Michigan Transportation Research Institute
2901 Baxter Road
Ann Arbor, MI 48152
umtri-ccat@umich.edu
(734) 763-2498

UMTRI | The University of Akron | Central State University | Illinois Center for Transportation | University of Minnesota | Northwestern | Purdue University | Washtenaw Community College | Wisconsin – University of Wisconsin–Madison

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

Technical Report Documentation Page

1. Report No. CTS 26-06	**2. Government Accession No.**	**3. Recipient's Catalog No.**
4. Title and Subtitle Cybersecurity of next-generation traffic signal control with CAVs		**5. Report Date** March 2026 **6. Performing Organization Code** Enter any/all unique numbers assigned to the performing organization, if applicable.
7. Author(s) Raphael Stern, PhD, 0000-0001-6633-7827 Amirhossein Kiani, 0009-0004-9047-0352		**8. Performing Organization Report No.** Enter any/all unique alphanumeric report numbers assigned by the performing organization, if applicable.
9. Performing Organization Name and Address Center for Connected and Automated Transportation University of Michigan Transportation Research Institute 2901 Baxter Road Ann Arbor, MI 48109 University of Minnesota 200 Oak Street SE Minneapolis, MN 55455		**10. Work Unit No.** **11. Contract or Grant No.** Contract No. 69A3552348305
12. Sponsoring Agency Name and Address U.S. Department of Transportation Office of the Assistant Secretary for Research and Technology 1200 New Jersey Avenue, SE Washington, DC 20590		**13. Type of Report and Period Covered** Final Report (September 2023 – December 2025) **14. Sponsoring Agency Code** OST-R
15. Supplementary Notes Conducted under the U.S. DOT Office of the Assistant Secretary for Research and Technology's (OST-R) University Transportation Centers (UTC) program.		
16. Abstract This report explores cyber vulnerabilities that exist in new approaches for actuated traffic flow management where individual connected and automated vehicles share information about their arrival to a reinforcement learning based traffic signal. Using this approach, potential cyberattacks include false data injection attacks, where vehicles report false data. We find that new signal techniques are vulnerable to such attacks, deteriorating intersection level of service, and also present approaches to overcome this vulnerability through federation.		
17. Key Words Traffic signal timing; Connected and automated vehicles; Transportation cybersecurity	**18. Distribution Statement** No restrictions.	

UMTRI · The University of Akron · Central State University · Illinois Center for Transportation · University of Minnesota · Northwestern · Purdue University · Washtenaw Community College · Wisconsin – University of Wisconsin–Madison

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	19	

Table of Contents

List of Figures

Introduction

Connected and automated vehicles (CAVs) provide new opportunities for improved traffic management. One area where CAVs may improve traffic flow management is in traffic signal control, where connectivity can be used to convey anticipated vehicle arrivals at an intersection, and this information can be used to develop improved traffic signal controllers. One such traffic signal control algorithm that has recently shown promise is the use of reinforcement learning (RL) to learn an adaptive traffic signal control. However, since these RL-based controllers depend on vehicle counts that are conveyed via connectivity from individual CAVs on the road, they provide a potential attack surface, where an attacker may inject false data through the communication network.

As traffic signals become increasingly connected through intelligent transportation systems, their vulnerability to cyberattacks is a growing concern. Cyberattacks on traffic networks can disrupt road infrastructure, causing severe economic and societal impacts. For example, in 2021, a cyberattack on Iran's transport ministry delayed and canceled hundreds of trains, and in 2017, attackers deactivated traffic cameras in Washington, DC. Researchers have also identified vulnerabilities in over 40 US cities *Traffic light hackers can cause road chaos* (n.d.); *Full Stop: Vulnerabilities in IoT Traffic Light Systems* (n.d.), showing that actuated traffic signals can be spoofed by manipulating sensor data. Similarly, connected vehicle systems are susceptible to Sybil attacks in V2I communication *Vulnerability analysis and defense framework for the cybersecurity of a traffic control system*.

These incidents highlight the critical need for cybersecurity in transportation systems as traffic sensors, intelligent controls, and connected vehicles become widespread. Studies like Maiti & Chilukuri (2021) demonstrate the effectiveness of reinforcement learning (RL) algorithms, such as Q-Learning, in optimizing signal timing and reducing delays. Likewise, Prashanth & Bhatnagar (2010) addresses scalability in adaptive traffic control, achieving efficiency in large-scale networks. This study examines the vulnerabilities of RL-based traffic signal controllers to cyberattacks. While traditional actuated controllers offer analytical bounds under data deterioration, RL-based controllers lack such guarantees. We analyze the impact of cyberattacks on RL-based controllers under various scenarios to assess their performance degradation.

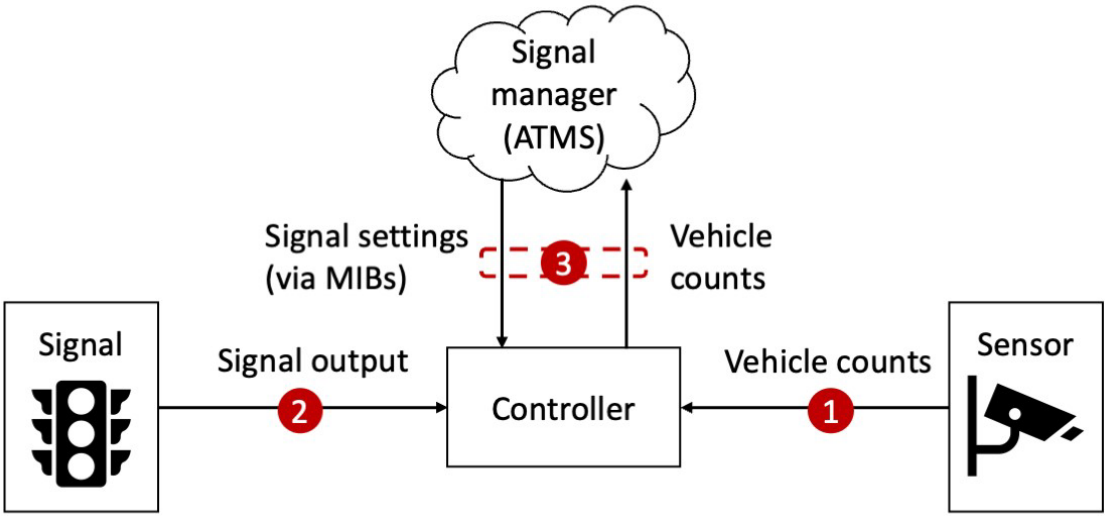

Figure 1: Traffic signal network architecture with possible attack locations numbered in red circles.

Background and Review of Literature

In this section, we review the state of the literature on actuated traffic signal control and traffic signal security.

Actuated Traffic Signal Control

Actuated traffic signal controllers manage signal phases based on vehicle arrivals detected by sensors. These systems, increasingly connected via communication networks, leverage real-time data for efficiency Feng et al. (2018); Wang et al. (2017). Adaptive algorithms like SCOOT Hunt et al. (1981) and SCATS Sims (1979) adjust signal timing based on real-time vehicle counts Stevanovic et al. (2009), while methods like Max Pressure Varaiya (2013) optimize throughput using traffic pressure. Despite their utility, these controllers rely on static assumptions about traffic flow, limiting adaptability to dynamic conditions Wei et al. (2021).

Reinforcement Learning in Traffic Control

Reinforcement learning (RL) has been widely studied for adaptive traffic signal control Abdulhai et al. (2003); Yau et al. (2017). RL agents adjust signal timing based on states (e.g., queue lengths) to minimize delays Mannion et al. (2016). RL approaches optimize phase selection Wei et al. (2018); Nishi et al. (2018) or cycle duration Aslani et al. (2017), outperforming traditional methods in various scenarios. Despite advances, these techniques lack robust defenses against cyberattacks, making them vulnerable to disruptions Feng et al. (2018).

Traffic Signal Cyberattacks

Increased connectivity introduces vulnerabilities to cyberattacks, such as sensor spoofing, Stuxnetlike attacks, and man-in-the-middle (MITM) attacks Vinayaga-Sureshkanth et al. (2020); Ezell et al. (2013). These attacks manipulate sensor data, controller logic, or management messages, leading to malicious signal control and traffic delays. This study focuses on the impact of spoofing attacks on traffic networks.

Contributions

This paper evaluates the impact of data integrity attacks on traffic networks, highlighting vulnerabilities in RL-based adaptive signal controllers. While these controllers outperform fixed-time methods under normal conditions, their performance degrades significantly under cyberattacks. Using real-world intersections and authentic data, this study emphasizes the need for resilient RL-based controllers to enhance traffic management reliability.

Findings

System setup

RL traffic signal controller

In this paper, we focus on the RL-based traffic signal control problem where RL agents control the phase changing of the traffic light based on the traffic condition at the intersections locally.

Let $N = \{1,\ldots,N\}$ be the set of intersections in the traffic network. The Markov decision process (MDP) associated with each RL agent $i \in N$ is defined as follows:

1. *State Space:* At each time step t, the state is defined as $s^i_t = (e,d)^i_t$, where e and d are the vectors of the arrival and departure counts for all lane groups at intersection i. Then, system state at time t is $s^i_t \in S_i$, where S_i is a discrete state space at the intersection i.
2. *Action Space:* Observing s^i_t, the agent will predict next action a^i_t that changes a certain phase in the intersection i. Due to the non-homogeneity of the intersections, the action space, denoted by A_i, is the set of all possible phases available for the intersection $^{i \text{ at } t}_{N}$, i.e., $a^i_t \in A_i$. Therefore, the next actions are defined as $a_t = (a^1_t,\ldots,a^N_t)^T \in A$, where $A = \times_{i=1} A_i$.

3. *Reward Function:* The actions will be executed in the environment, and the rewards $r_t = (r_t^1 \ldots, r_t^N)^T$ will be generated. The local reward function is defined as $r_t^i = r\left(s_t^i, a_t^i\right)$ rep-

resenting how much the sum of the waiting times (waiting time refers to the total amount of time that vehicles are stationary at a red light.) of all approaching vehicles has changed between successive actions in the intersection i. In other words, let $D_{a_t^i}$ be the sum of the waiting times of all approaching vehicles in the intersection i after the execution of an action a_t^i ; then, $r_t^i = D_{a_{t-1}^i} - D_{a_t^i}$. A vehicle is classified as waiting if its speed falls below 0.1 m/s. Based on this definition, a greater reduction in cumulative waiting time results in a higher reward. Thus, by maximizing these rewards, the agents work to minimize waiting times at intersections, consequently enhancing the local traffic flow.

For each intersection i at time t, let $\pi_t{}^i : S_i \rightarrow P(A_i)$ denote a policy that maps a state $s_t{}^i$ to a probability distribution over the action space A_i where $P(A_i)$ is the set of probability distributions over the action space A_i. For each intersection i, we seek a phase-changing policy $\pi_i = (\pi_1^i, \pi_2^i, \ldots)$,

such that $a^i{}_t \sim \pi_t{}^i(a_t|s^i{}_t)$, $a^i{}_t \in A_i$. Starting from an initial state s_0^i, the i^{th} traffic signal controller seeks a policy $\pi_i{}^* \in \Pi_i$ that maximizes total expected rewards, i.e.,

$$\pi_i^* \in \arg\max_{\pi \in \Pi} \mathbb{E}\left(\sum_{t=1}^{T} r_t^i \middle| \pi_i, \boldsymbol{s}_0^i\right) \tag{1}$$

where Π_i is the set of all feasible phase-changing policies.

Solution Approach

Solving the MDP in Eq. (1) is challenging due to the lack of perfect environmental knowledge in traffic signal control and the computational infeasibility of traditional Q-learning for large state spaces. Attacks further complicate the problem by altering system states, rendering tabular Qlearning inadequate as it cannot handle unseen state-action pairs effectively.

Function approximation addresses the limitations of Q-learning by estimating Q-values for stateaction pairs, enabling scalability in large state spaces Maiti & Chilukuri (2021); Prashanth & Bhatnagar (2010). However, traditional function approximation methods require careful selection of functional forms, risking either high bias with overly simple models or high variance with overly complex ones. Balancing this trade-off is nontrivial and requires extensive tuning.

To address these challenges, we employ Deep Q-Networks (DQN), which approximate the Qfunction using neural networks. DQNs are scalable and can handle large state spaces by generalizing across similar states, allowing effective responses to new or altered states, including those affected by attacks. Unlike traditional methods, DQNs eliminate the need for predefined functional forms, offering greater flexibility and adaptability for real-time traffic signal control.

Attack model for traffic signal controls

In this paper, we focus on the sensor spoofing attack that tempers with the data from the vehicle detectors and injects faulty vehicle count data to trick the RL-based traffic controllers into adopting compromised or even dangerous traffic signal phasing Laszka et al. (2016). This type of practical physical attack does not require direct access to the signal controller; however, they seek physical consequences by interfering with and manipulating the sensor data.

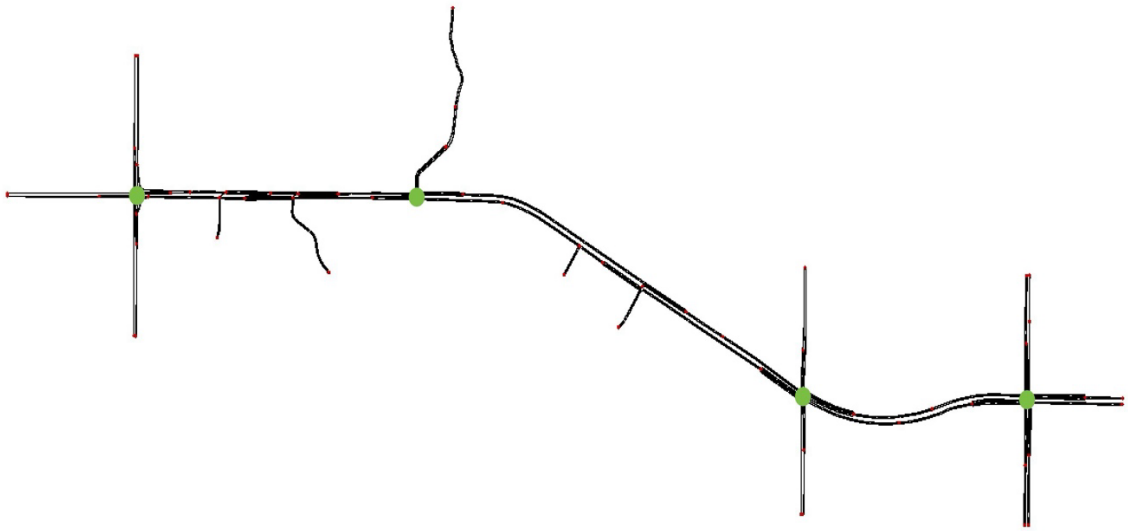

Figure 2: Simulation network in Hennepin County, Minnesota with four intersections as shown with the green circles. Each of the four intersections is controlled by an RL-based controller in simulation and real-world demand patterns are provided by Hennepin County for both peak and off-peak traffic flow conditions.

Each RL-based traffic signal controller makes phase-changing decisions, $a^i{}_t$, based on their realtime observation of the environment, $s^i{}_t$, that originates from the detectors such as loop detectors, CCTV, etc. In this context, the adversary first injects fake data into the sensor data communication channels. These malicious data are received by the RL controller as the agent's observation, causing them to change the phase of the traffic light incorrectly. The attack can be modeled as

$$s_t^{i'} = s_t^i + \alpha_t^i \quad (2)$$

where α^i_t represents the bias caused by false data injection by the attacker to the original sensor readings, s^i_t, to trick the RL-based traffic signal controller. Accordingly, the RL-based traffic signal controller makes phase-changing decisions based on the manipulated observation $s^i_{t'}$. Such false perceptions can adversely affect an agent's decision-making in subsequent stages of signal phasing. Incorrect timing of traffic signals might not streamline traffic flow, leading to unnecessary idling at intersections or circulating in search of alternative routes. Such ineffective traffic management leads to increased congestion. Moreover, it increases the likelihood of accidents due to drivers' irritation and unpredictable driving as they deal with slow or stopped traffic conditions.

Simulation study

Next, we present a simulation study where we demonstrate the ability of adaptive RL-based controllers to outperform traditional fixed-time controllers and the potential for RL-based adaptive controllers to have deteriorating performance under data-integrity attacks. First, we introduce the simulation setup, and then we present the simulation results.

Simulation environment

SUMO, an open-source-based traffic simulation software, initially released in 2001, has enabled transportation researchers to simulate numerous scenarios. When combined with RL, SUMO provides interfaces that can be used to evaluate the efficiency of the model. One of the key features of SUMO is its development of the library TraCI Wegener et al. (2008) that provides users with a versatile toolset and access to the simulation. This library serves as a bridge between external applications and the simulation environment in SUMO. In this study, we implement the simulation using TraCI library in conjunction with Python 3.11. We study two main cases, one for peak hour and the other for off-peak hours traffic flow, provided by Hennepin County, MN. The network consists of four four-leg intersections, which have been depicted in the following figure:

DQN Setup

In this section, we provide a comprehensive explanation of the DQN architecture used in our study. Our implementation follows the architecture proposed in Mnih et al. (2015), designed to approximate the Q-function through a neural network structure. We considered a scenario where all traffic lights are controlled by a reinforcement learning agent, with each agent have the same neural network structure and work independently. The neural network consists of an input layer, a hidden layer, and an output layer.

The input to the neural network is the state representation, which includes the vectors of arrival and departure counts for all lane groups at intersection i, denoted as $s^i_t = (e,d)^i_t$. This layer is followed by a SeLU (Scaled Exponential Linear Unit) activation function and batch normalization to enhance the stability and performance of the network. The hidden layer is a fully connected layer with 100 neurons, also followed by a SeLU activation function and batch normalization. The purpose of this layer is to transform the input state into a higher-dimensional feature space, facilitating the learning of complex representations. The output layer is a fully connected layer with $|A_i|$ neurons, where $|A_i|$ is the number of possible phases available for intersection i. Each neuron in this layer corresponds to the probability of a specific action given the input state, and this layer is followed by a softmax activation function, which converts the outputs into a probability distribution over the possible actions.

For training, we use a replay buffer of size $100{,}000$ to store past experiences. During training, mini-batches of experiences are sampled from the replay buffer to break temporal correlations and stabilize training. A separate target network is maintained to provide stable target probabilities. The target network is periodically updated with the weights of the online Q-network, with an update frequency of 10 steps and an update rate of $\tau = 0.05$. The network is optimized using the RMSprop optimizer with a learning rate of 0.01. The discount factor γ is set to 0.99, balancing immediate and future rewards. An ϵ–greedy approach is applied for exploration, with ϵ annealed from 0.95 to 0.05 over the first 2700 steps (equivalent to 45 minutes). During the learning phase, the environment is simulated for 10 episodes, each consisting of 200 steps. Each step corresponds to 3 time units of simulation, allowing the agent to learn and adapt its policy over multiple scenarios and time frames.

Numerical setup

For this study, we consider an attack on all the intersections in the network depicted in Fig 2. The attacks we design for the purpose of this study change the input state of the RL model. In this regard, we define the variable $\alpha^i_t \sim N(\mu_\alpha, 1)$ that represents the amount of added noise to the corresponding agent's model. As mentioned before, this noise comes in the form of the number of vehicles on the incoming and outgoing lanes of the intersection. After that a noise is randomly choosen, this noise is added to the state space of the reinforcement learning model of one of the phases that is also randomly picked. In order to control the one-sided impact of the attack, the absolute value $\|\alpha_t\|$ is added to the corresponding agent's state, s^i_t. The mean μ_α in this case represents the severity of the

attack. We run the experiments at six levels of attack severity, where $\mu_\alpha \in \{1,2,3,4,5\}$ with 1 and 5 being the least severe and most severe attacks. These values are chosen based on the average number of vehicles arriving at each intersection. We avoided having negative values for the noise, as it might cancel out the impact of noises that has been added to the agents' state space previously. For each μ_α level as well as for the scenario with no attack, 10 replications are implemented to obtain the average performance and the confidence interval. To depict the distribution of number of vehicles, we use figure 3. This figure depicts the number of active vehicles in each of peak and off-peak hour scenarios. Active vehicle is defined as a vehicle that has entered the network from the arrival lanes but has yet to leave the network.

Analysis of Simulation Results

The simulation results for the total number of stopped vehicles under off-peak and peak-hour conditions are presented in Figures 4 and 5, respectively. Both scenarios simulate 45 minutes of traffic data, using the total number of stopped vehicles across the signal corridor as the performance metric.

The baseline fixed-time controller, based on Webster's method, is shown in black and performs consistently across repetitions due to its deterministic nature. In contrast, the RL-based adaptive

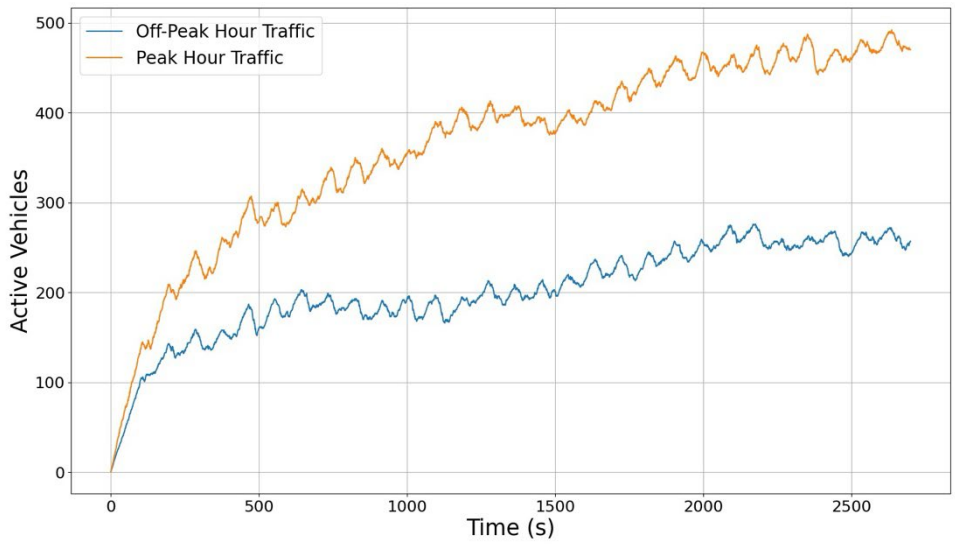

Figure 3: Number of active vehicles in peak and off-peak hour scenarios.

controller, under typical (unattacked) conditions (blue line), outperforms the fixed-time controller in both off-peak and peak-hour scenarios by dynamically adjusting signal timings to traffic demand. However, as data integrity attacks intensify (green, red, purple, orange, and cyan lines), the RL controller's performance deteriorates, while the fixed-time controller remains unaffected.

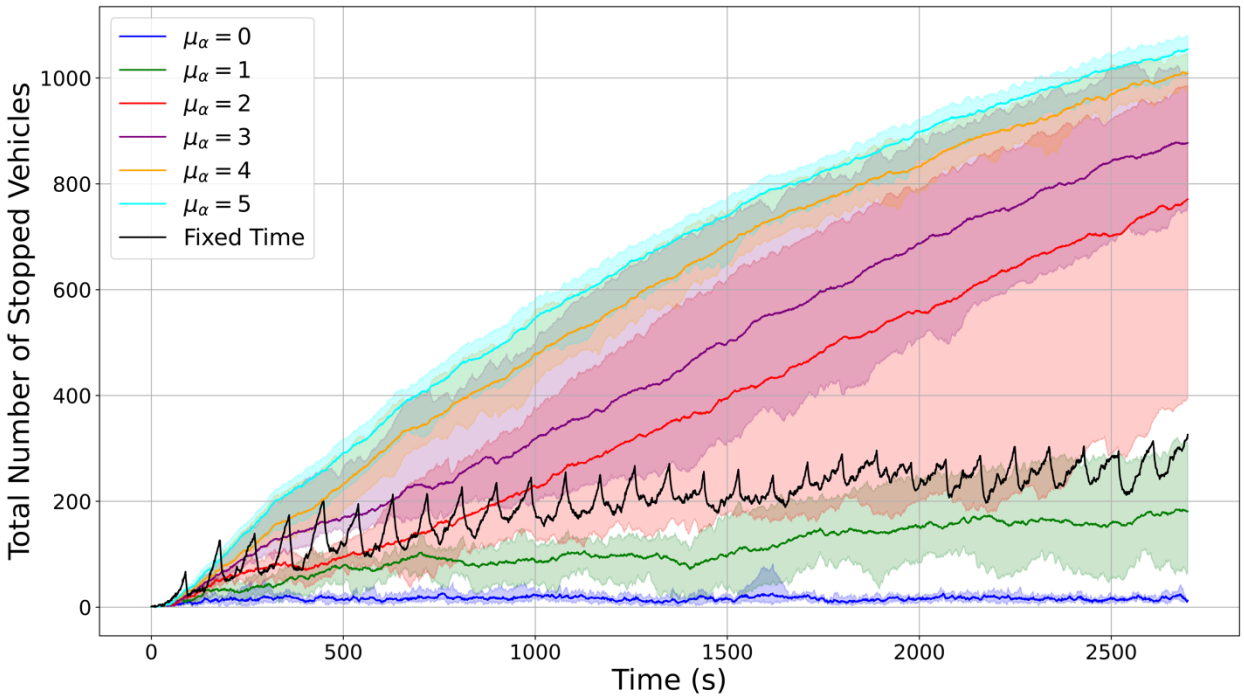

Figure 4: Off-peak hour results showing total stopped vehicles over time for RL-based and fixed-time controllers. RL controller performance deteriorates under increased attack intensity.

During off-peak conditions (Figure 4), the RL-based controller reduces stopped vehicles compared to the fixed-time controller under unattacked conditions but performs worse as attack intensity increases. In peak-hour conditions (Figure 5), the incoming demand is so high that all controllers become unstable. Still, the RL-based controller generally performs better than the fixed-time controller under unattacked conditions. Under data integrity attacks, however, RL-based controller performance degrades significantly, highlighting its reliance on accurate data.

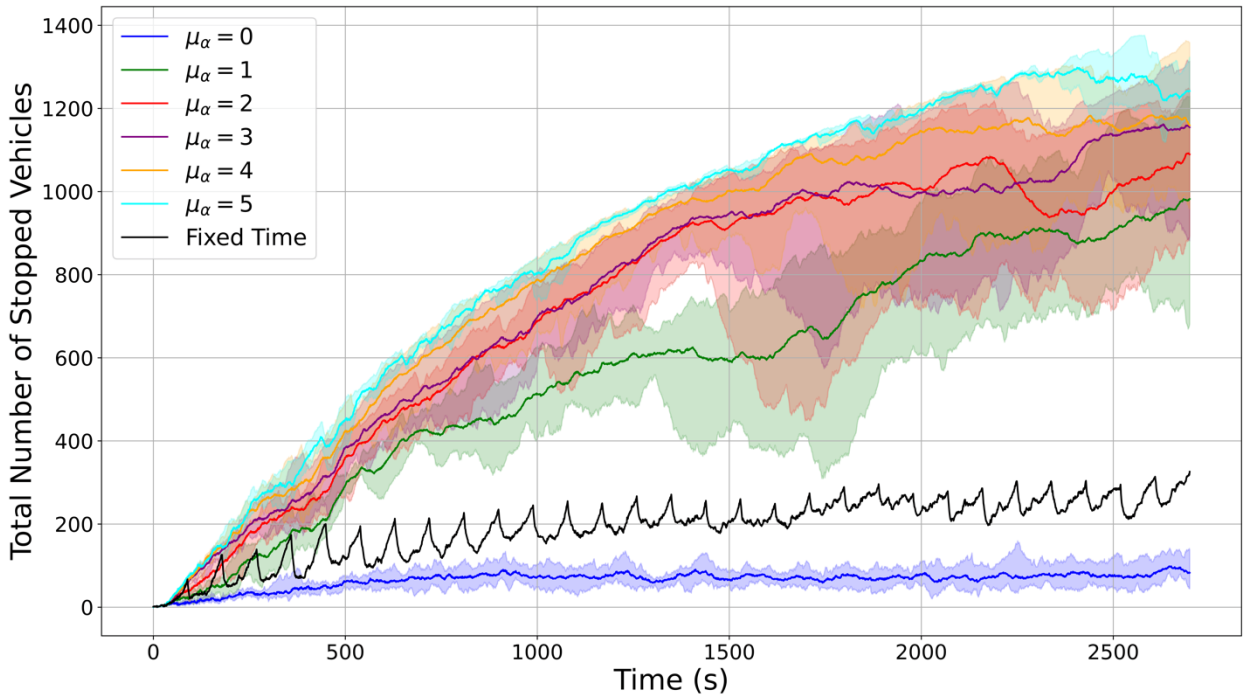

Figure 5: Peak hour results showing total stopped vehicles over time for RL-based and fixed-time controllers. While all controllers are unstable, RL-based controllers outperform fixed-time controllers under unattacked conditions but deteriorate under attacks.

Reward Aggregation Mechanism

The reward aggregation mechanism was optimized to improve the scalability and performance of RL-based controllers. Initially, a global reward function aggregated inputs from all agents, but it proved inefficient due to noise from distant intersections. To address this, a localized mechanism was introduced, focusing on rewards from an agent and its immediate neighbors:

$$R_t^g(S_t) = f_{agg}(R_t^i(s_t^i), R_t^{adj}(s_t^{adj}); \theta^i)$$

Experimental results (Figure 6) show that the localized mechanism reduces the total number of stopped vehicles and improves scalability, particularly under high traffic densities.

Key Findings:

- Scalability: Localized aggregation consistently outperforms global aggregation under higher traffic densities.
- Reduced Stopping Events: Localized aggregation significantly reduces stopped vehicles, especially when noise is minimized.

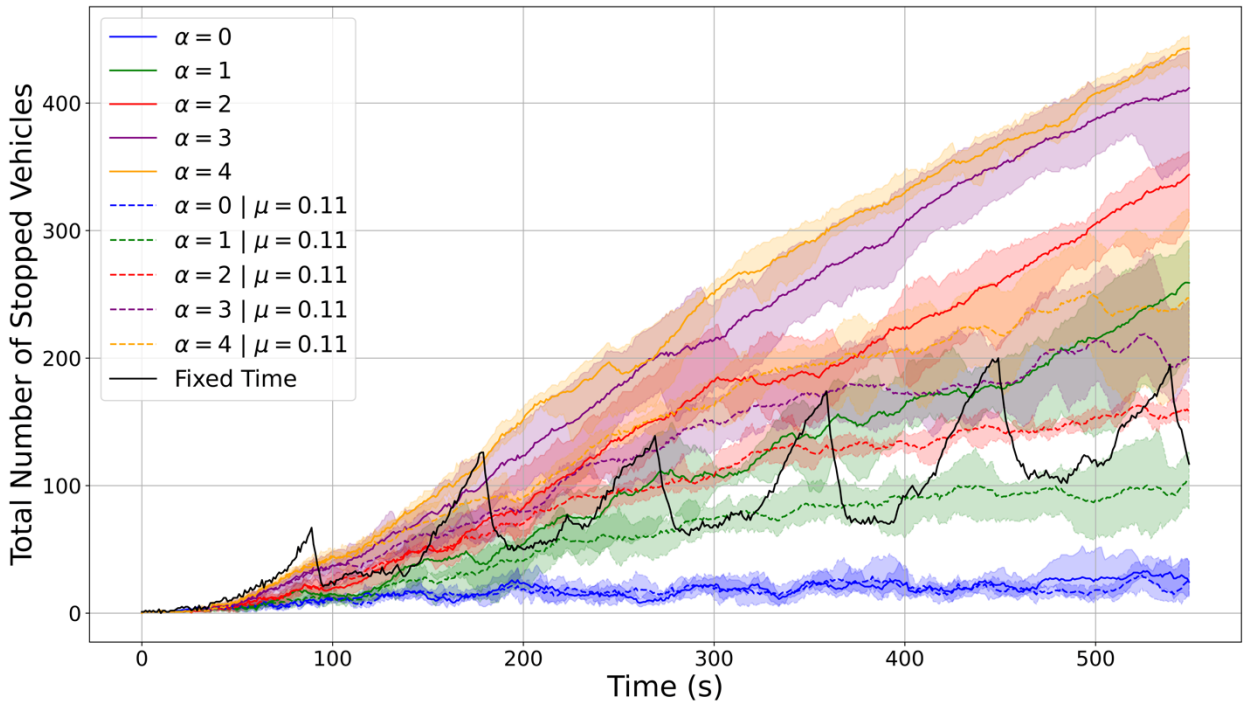

Figure 6: Performance under μ = 0.11. Localized aggregation reduces stopped vehicles compared to global aggregation.

Recommendations

This study evaluates the impact of data integrity attacks on RL-based actuated traffic signals in road transportation networks. More specifically, through several simulated attacks applied to the real traffic network, it has been shown that data integrity attacks that spoof the sensor data can significantly increase the traffic delay at the intersections controlled by RL-based actuated signals. The results indicate the importance of cybersecurity for signalized traffic networks to avoid potential severe adverse economic and societal impacts that cyberattacks can cause if they target road traffic infrastructure systems. This work is timely as AI, especially RL-based control, and IoT technology are getting closer attention and broader applications in transportation applications with more internet-based digital communications. The work highlights the need for a better-secured smart signal control framework in terms of not only the network protocols but also the learning and control mechanisms. This work lays the foundation for future studies in this under-explored area, which includes investigating the ripple effects of data-integrity attacks on more complex and larger-scale traffic networks as well as designing cyberattack-resilient actuated signals that can adapt their policy to mitigate the impact of an attack on the traffic.

Final Documentation of Outputs

Synopsys of Performance Indicators

Outputs description

- **Presentations:**
    - Poster presentation: " Quantifying The Impacts of Data Integrity Attacks on RL-Based Traffic Signal Control," Transportation Research Board Annual Meeting, Washington, DC, January 6, 2025.
    - Conference talk: "Cybersecurity of Connected and Automated Vehicles via Traffic Anomaly Detection," CCAT Global Symposium, Ann Arbor, MI, March 28, 2025.
    - Conference plenary address: "Leveraging vehicle automation to improve traffic conditions for all road users," IFAC International Symposium on Advances in Automotive Control, Eindhoven, Netherlands, June 18, 2025.
    - Departmental seminar: "Modeling and control of heterogeneous and mixed-autonomy traffic flow," Department of Civil, Environmental, and Geo-Engineering Warren Lecture, University of Minnesota, Minneapolis, MN, September 13, 2024.
- **Increase in the body of knowledge:**
    - Provided an understanding of cyber vulnerabilities that arise in connected and automated vehicle-based traffic signal control.
    - Provided an understanding of possible attack vectors for cyberattacks on connected and automated vehicles.
- **Enlargement in the pool of trained transportation professionals:**
    - Two graduate students received funding on this project. One is now an assistant professor at an R1 research university, and the other has applied the knowledge gained in this project to an internship with a local county department of public works.

Outcomes and Impacts description

- **Increase in the understanding and awareness of transportation issues:**
    - The funded research has been shared with the academic community, professionals, and the general public.

    - The research project shared awareness on both opportunities for improved traffic management with connected and automated vehicles through next-generation signal control that leverages connectivity, as well as risks for compromise of these control techniques.
- **Increase in the body of knowledge:**
    - The research project increased the body of knowledge on vehicle connectivity, automation, and how to use these features to improve traffic management.
    - The research project increased the body of knowledge on risks associated with cyberattacks that compromise communication networks for connected and automated vehicles.

Challenges and lessons learned

Challenges:

- Connectivity and automation are still somewhat speculative technologies. While prototypes exist, widespread adoption is not prevalent, so any modeling work is based on assumptions of how these technologies will be adopted, instead of observations of current adoption patterns. This means any results are only as valid as the assumptions they are based on, and presented a challenge for the research team.
- While vehicle behavior can easily be modeled, human behavior is far more difficult to model, and more probabilistic. Therefore, simplifying assumptions on human behavior and compliance are made, but predicting true system behavior presented a challenge for the research project.

Lessons learned:

- Transportation systems may greatly benefit from connectivity and automation through improved operational efficiency (e.g., improved traffic signal timing, etc.).
- These next-generation traffic management techniques are susceptible to compromise through cyberattacks on the communication layer.

Works Cited

Abdulhai, B., Pringle, R., & Karakoulas, G. J. (2003). Reinforcement learning for true adaptive traffic signal control. Journal of Transportation Engineering, 129(3).

Aslani, M., Mesgari, M. S., & Wiering, M. (2017). Adaptive traffic signal control with actor-critic methods in a real-world traffic network with different traffic disruption events. Transportation Research Part C: Emerging Technologies, 85, 732–752.

Ezell, B. C., Michael Robinson, R., Foytik, P., Jordan, C., & Flanagan, D. (2013). Cyber risk to transportation, industrial control systems, and traffic signal controllers. Environment Systems and Decisions, 33, 508–516.

Feng, Y., Huang, S., Chen, Q. A., Liu, H. X., & Mao, Z. M. (2018). Vulnerability of traffic control system under cyberattacks with falsified data. Transportation research record, 2672(1), 1–11.

Full stop: Vulnerabilities in IOT traffic light systems. (n.d.). https://www.bankinfosecurity.com/ full-stop-vulnerabilities-in-iot-traffic-light-systems-a-14945. (Accessed: 202203-21)

Hunt, P., Robertson, D., Bretherton, R., & Winton, R. (1981). Scoot-a traffic responsive method of coordinating signals (Tech. Rep.).

Laszka, A., Potteiger, B., Vorobeychik, Y., Amin, S., & Koutsoukos, X. (2016). Vulnerability of transportation networks to traffic-signal tampering. In 2016 acm/ieee 7th international conference on cyber-physical systems (ICCPS) (pp. 1–10).

Maiti, N., & Chilukuri, B. R. (2021). Traffic signal control for an isolated intersection using reinforcement learning. In 2021 international conference on communication systems & networks (COMSNETS) (pp. 629–633).

Mannion, P., Duggan, J., & Howley, E. (2016). An experimental review of reinforcement learning algorithms for adaptive traffic signal control. Autonomic road transport support systems, 47–66.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Ostrovski, G. (2015). Human-level control through deep reinforcement learning. nature, 518(7540), 529–533. Nishi, T., Otaki, K., Hayakawa, K., & Yoshimura, T. (2018). Traffic signal control based on reinforcement learning with graph convolutional neural nets. In 2018 21st International Conference on Intelligent Transportation Systems (ITSC) (pp. 877–883).

Prashanth, L., & Bhatnagar, S. (2010). Reinforcement learning with function approximation for traffic signal control. IEEE Transactions on Intelligent Transportation Systems, 12(2), 412–421.

Sims, A. G. (1979). The Sydney coordinated adaptive traffic system. In Engineering foundation conference on research directions in computer control of urban traffic systems, 1979, pacific grove, California, USA.

Stevanovic, A., Kergaye, C., & Martin, P. T. (2009). Scoot and scats: A closer look into their operations. In 88th annual meeting of the transportation research board. Washington DC.

Traffic light hackers can cause road chaoshttps://www.wired.co.uk/article/hacking

-traffic-lights. (Accessed: 2022-03-21)

Varaiya, P. (2013). Max pressure control of a network of signalized intersections. Transportation Research Part C: Emerging Technologies, 36, 177–195.

Vinayaga-Sureshkanth, N., Wijewickrama, R., Maiti, A., & Jadliwala, M. (2020). Security and privacy challenges in upcoming intelligent urban micromobility transportation systems. In Proceedings of the second ACM workshop on automotive and aerial vehicle security (pp. 31–35).

Vulnerability analysis and defense framework for the cybersecurity of a traffic control system. (n.d.). https://www.roadsbridges.com/vulnerability-analysis-and-defense-framework

Wang, M., Winbjork, M., Zhang, Z., Blasco, R., Do, H., Sorrentino, S., ... Zang, Y. (2017). Comparison of lte and dsrc-based connectivity for intelligent transportation systems. In 2017 IEEE 85th vehicular technology conference (VTC Spring) (pp. 1–5).

Wegener, A., Piórkowski, M., Raya, M., Hellbrück, H., Fischer, S., & Hubaux, J.-P. (2008). Traci: an interface for coupling road traffic and network simulators. In Proceedings of the 11th communications and networking simulation symposium (pp. 155–163).

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

Wei, H., Zheng, G., Yao, H., & Li, Z. (2018). Intellilight: A reinforcement learning approach for intelligent traffic light control. In Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining (pp. 2496–2505).

Yau, K.-L. A., Qadir, J., Khoo, H. L., Ling, M. H., & Komisarczuk, P. (2017). A survey on reinforcement learning models and algorithms for traffic signal control. ACM Computing Surveys (CSUR), 50(3), 1–38.