CECS
CENTER FOR EMBEDDED & CYBER-PHYSICAL SYSTEMS

An Integrated Corridor Management for Connected Vehicles and Park and Ride Structures using Deep Reinforcement Learning

Yasamin Moghaddas* Mohamad Fakih*, Tyler Zhang, Mohanad Odema, and Mohammad Abdullah Al Faruque

* Both authors contributed equally to this work

Center for Embedded and Cyber-Physical Systems

University of California, Irvine

Irvine, CA 92697-2620, USA

{ymoghadd, mhfakih, tylerz1, modema, alfaruqu}@uci.edu

CECS Technical Report 23-03

August 16, 2023

An Integrated Corridor Management for Connected Vehicles and Park and Ride Structures using Deep Reinforcement Learning

Yasamin Moghaddas* Mohamad Fakih*, Tyler Zhang, Mohanad Odema, and Mohammad Abdullah Al Faruque
Department of Electrical Engineering and Computer Science, University of California, Irvine
{ymoghadd, mhfakih, tylerz1, modema, alfaruqu}@uci.edu

***Abstract*—The upcoming Connected Vehicles (CV) technology shows great promise in effectively managing traffic congestion and enhancing mobility for users along transportation corridors. Data analysis powered by sensors in CVs allows us to implement optimized traffic management strategies optimizing the efficiency of transportation infrastructure resources. In this study, we introduce a novel Integrated Corridor Management (ICM) methodology, which integrates underutilized Park-And-Ride (PAR) facilities into the global optimization strategy.**

To achieve this, we use vehicle-to-infrastructure (V2I) communication protocols, namely basic safety messages (BSM) and traveler information messages (TIM) to help gather downstream traffic information and share park and ride advisories with upstream traffic, respectively. Next, we develop a model that assesses potential delays experienced by vehicles in the corridor. Based on this model, we employ a novel centralized deep reinforcement learning (DRL) solution to control the timing and content of these messages. The ultimate goal is to maximize throughput, minimize carbon emissions, and reduce travel time effectively.

To evaluate our ICM strategy, we conduct simulations on a realistic model of Interstate 5 using the Veins simulation software. The DRL agent converges to a strategy that marginally improves throughput, travel speed, and freeway travel time, at the cost of a slightly higher carbon footprint.

I. Introduction

With the continuous growth of metropolitan cities, traffic congestion can become a concern. According to the latest Urban Mobility Report [1], the annual person-hours delay of 2020 was 46, resulting in a total of $1.14 billion for congestion cost. Freeways are known for experiencing a great deal of traffic congestion [2] as more people travel on the freeway simultaneously, resulting in slower speeds.

As a result, there have been works that focus on the notion of integrated corridor management (ICM) [3], which promotes the adoption of global traffic management strategies by integrating various traffic components along a corridor (e.g. ramp meter controllers) into a single interconnected system that aims to minimize congestion. Furthermore, connected vehicles (CV) can benefit from ICM through Dedicated Short-Range Communications (DSRC), vehicle-to-vehicle (V2V), and vehicle-to-infrastructure (V2I) standards [4]. ICM relies primarily on V2I communication instead of V2V. CVs allow for V2V communication, which is lower in latency because it is the exchange of information directly between vehicles. Having a sweeping view of surrounding vehicles offers more information about road conditions in addition to V2I communication. Therefore, this connectivity offers numerous advantages, such as transmitting traffic information to commuters, enabling priority requests for certain vehicles, and supporting features like platooning and adaptive cruise control [4].

The combined potential of ICM strategies and CVs could be enhanced using underutilized resources, such as park-and-ride (PAR) facilities. These facilities offer affordable and accessible public transport allowing individuals to park their vehicles safely and then travel to their destination via a local bus transit, carpool, etc [5]. PAR facilities are not being fully utilized in California, with most facilities only 65% full on average during peak hours [6, 7]. Integrating PAR supply and demand dynamics within ICM optimization strategies can improve traffic flow by encouraging greater usage of PAR structures.

This report presents a centralized learning-based approach that uses traffic information from the corridor and parking availability data from PAR facilities to formulate a corridor-wide advertisement strategy. The deep reinforcement-learning (DRL) approach processes the observed congestion state of the corridor and uses that information to guide the message advertisement strategies of the roadside units (RSUs) to vehicles traversing the corridor. The RSUs broadcast PAR advisory messages along the corridor. A DRL approach was specifically chosen because we have no initial knowledge about the traffic states of the corridor and corresponding message advertisement strategies that could be provided by a training dataset. A DRL model addresses this issue by gradually learning patterns that maximize a reward function that depends on throughput, travel time, and CO2 emissions. The goal is to increase PAR usage and efficiency and to improve corridor flow. The main contributions include:

- A realistic traffic model for a congested highway corridor that depicts the characteristics and roles of different traffic components
- Demonstration of how park-and-ride parking messages can be incorporated within DSRC
- A novel Deep Reinforcement Learning methodology that optimizes traffic flow and carbon emissions along a corridor
- Experiments using SUMO and VEINS, which simulate traffic conditions and our centralized ICM strategy on

* Both authors contributed equally to this work.

Interstate 5, indicating improvements in throughput, CO2 emissions, and travel time.

II. Background

A. DSRC Technology

Dedicated short-range communication (DSRC) is a wireless communication technology proposed by the United States Federal Communications Commission (FCC) that enables vehicle-to-everything (V2X) communication in intelligent transportation systems (ITS). The FCC reserved 75 MHz of spectrum from the 5.9 GHz frequency band for transportation-related communications [8]. DSRC uses IEEE 802.11p as its communication standard for its PHY and MAC layers. For the Security, Network and Transport, and Upper MAC layers implement IEE 1609.2, IEEE 1609.3, and IEEE 1609.4 respectively as can be seen in Table I.

Network Layer	Protocol
Application	IEEE 1609.1
Messaging Sublayer	SAE J2735 (V2X), SAE J2945.1 (V2V)
Security	IEEE 1609.2
Network and Transport	IEEE 1609.3
Upper MAC	IEEE 1609.4
Lower MAC	IEEE 802.11p
Physical	IEEE 802.11p

TABLE I: Network layers and their corresponding protocols

DSRC packets contain data related to a vehicle's speed and position, for example. In traffic management applications the packets could contain data relevant to traffic congestion or road conditions. The system itself consists of On-Board Units (OBUs) and Roadside Units (RSUs). An OBU is a transceiver that is mounted on a vehicle, while an RSU is a transceiver that is mounted on the road. The OBUs in various vehicles exchange information with one another. RSUs communicate with OBUs and send information back to a centralized hub [9]. Our work focuses on leveraging V2I communications between OBUs and RSUs for our ICM strategy.

B. Related Work

1) Connected Vehicle Platoons: One common line of study in CV research is to use V2V coordination to organize traveling vehicles into groups or platoons [10, 11], but these tend to focus on local algorithms and lack system-wide awareness. One such study explored the benefits of utilizing three mobility improvements called cooperative adaptive cruise control, speed harmonization, and queue warning in a highway scenario [11]. The study concluded that individual vehicle safety was improved but at the cost of sacrificing overall highway throughput.

2) Ramp Metering: A popular system-wide ICM approach is to adjust ramp metering rates to control the flow of vehicles entering the highway [12, 13, 14, 15]. Several ramp metering algorithms such as ALINEA and SWARM have existed and have been in use since the early 2000s [12]. SWARM in particular can be found in Orange County; it calculates one ramp metering rate based on local density and another based on a predetermined global volume reduction, picking the most restrictive metering rate from the two [12].

Many advancements in ramp-metering approaches have been proposed since then. In 2014, Fares and Gomaa implemented a deep reinforcement algorithm to adjust ramp metering timings based on vehicle density that increased highway flow and reduced total travel time [13]. In another study in 2016, Hashemi and Abdelghany proposed a system that dynamically searches for an optimal traffic control scheme involving intersection timing plans, diversion messages, ramp metering, and/or dynamic shoulder lanes [14]. They demonstrated improvements to total travel time when more control options are available. In a follow-up work, Hashemi and Abdelghany trained a deep convolutional neural network to recognize the highway state and pick an optimal traffic control scheme [15]. These strategies do not explicitly mention CV technology, but the state vectors from Fares and Goma [13] and Hashemi and Abdelghany [15] could potentially be obtained through RSUs sampling data from nearby CVs.

An even more recent approach [16] involves a physics-informed reinforcement learning-based strategy for coordinated ramp metering. While the results are promising for total time spent savings, the authors acknowledge that a larger traffic network size should have been considered.

3) Dynamic Rerouting: Ramp metering seems to be a well-explored approach, so the scope of our work focuses on improving corridor traffic using dynamic rerouting to under-utilized PAR structures. Liu et al. [17] set out to establish a framework for evaluating ICM methods and conducted a study on diverting upcoming traffic to side roads using variable message signs on the freeway, using throughput and travel time as evaluation criteria. They demonstrate that diverting traffic reduces travel delay for freeway vehicles but degrades side road performance, so careful evaluation of trade-offs is required. Our ICM approach is similar to this study, so it makes sense to use highway throughput and vehicle travel time as our evaluation criteria as well.

A study conducted by Ortega et al. [18] demonstrated via simulation that use of a park and ride is expected to increase the total trip time compared to not using the resource. [19] also note that park and ride increased total trip time presumably because the mode of transportation for PAR is slower and the diverted routes taken to arrive at a PAR increase delay. This indicates that in our evaluation metrics, we must incorporate criteria beyond travel time, such as emissions or highway throughput.

4) Deep Reinforcement Learning in ICM Strategies: The studies conducted by Fares and Goma [13] and Hashemi and Abdelghany [15] utilize deep reinforcement learning to estimate the state of the network with relative success. In addition, Liu et al. [20] established a Markov decision process (MDP) framework for a freeway scenario. These examples serve as a foundation for establishing an MDP for our scenario and using deep learning to optimize the ICM strategy. However, our deep learning agent will differ from these studies in its action space of diverting traffic to PAR systems.

III. Scenario

A. Interstate 5

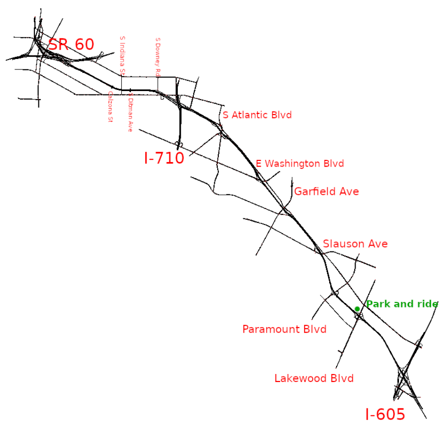

Fig. 1: Interstate 5 model

We conducted our study on Interstate 5, specifically the northbound 10 mile stretch between Interstate 605 and California State Road 60 in LA County. Interstate 5 is known for being one of the most congested corridors in California [21] and the United States [22, 23]. We used SUMO's OSMWebWizard tool to convert an OpenStreetMap rendering of I-5 into a SUMO road network 1. For this study, we considered vehicles traveling northbound from I-605, the source, to SR 60, the sink.

1) Realistic Traffic Modeling: The corridor includes a diverse set of on-ramps and exits with regard to traffic volume. Ramp volumes and average daily traffic data were collected from the Caltrans Traffic Census Program's [24] 2015 survey to generate realistic traffic flows.

Vehicles were spawned at each on-ramp as a Poisson process. The ramp volumes were recorded as daily averages [24]; however, we converted them to Poisson rate with units $\frac{vehicles}{second}$ as seen in equation 1.

$$\lambda_i = \frac{V_i \frac{veh}{day}}{24 * 60 * 60 \frac{sec}{day}} = \frac{V_i}{86400} veh/sec \tag{1}$$

where λ_i is the Poisson rate for on ramp i and V_i is its average daily volume.

To model the exit behavior, each vehicle is assigned a destination following a fixed probability distribution when it spawns (referred to as a route distribution [25]). We computed the probability of taking an exit according to equation 2.

$$\mathbb{P}_i = \frac{L_i}{\sum_{n=1}^{i} V_n} \tag{2}$$

where $\mathbb{P}_i$ is the probability of taking exit i, L_i is the average daily volume of exit i, and V_n is the average daily volumes for on ramps before exit i. Each on ramp has a unique route distribution table populated with $\mathbb{P}_i$ for the exits that are accessible downstream. If none of the exits are taken, the vehicle will route to the sink. For example, a vehicle at the on ramp of S Downey Rd (see figure 1) will have:

1) Chance to exit at Ditman Ave = $\mathbb{P}_{DitmanAve}$
2) Chance to exit at Calzona St = $\mathbb{P}_{CalzonaSt}$
3) Chance to exit at the sink = $1.0 - (\mathbb{P}_{DitmanAve} + \mathbb{P}_{CalzonaSt}) = \mathbb{P}_{sink}$

NOTE: Indiana St does not have an exit, so it is omitted.

2) Park and Ride: The North Lakewood Park and Ride facility is located near Lakewood Blvd. For this study, we assumed the availability of a bus schedule, and these buses spawn at a fixed rate in the simulation. The buses proceed directly to the highway exit without any detours to other exits.

B. Preliminary Exploration - The Addition of a Park-and-Ride Structure

From the simulation we noticed that diverting traffic to the North Lakewood Park and Ride facility leads to improvements in mainline flow, reduced CO2 emissions and overall speed of traffic. However, this improvement came at the cost of longer travel times for exiting vehicles. We started by using SUMO's traffic scaling parameter to adjust the spawn rates of the vehicles until congestion became evident, which happened when the spawn rates were doubled. We found that congestion was heaviest at Interstate 605 and Interstate 710 junctions due merging traffic.

In this scenario, we divert vehicles to the park and ride using three different exits at a fixed probability, which we call the compliance parameter. If the vehicle is compliant with being rerouted, then it has an equal chance of taking any of the three exits, namely I610, Lakewood Boulevard, and Paramount Boulevard. For this scenario we set the park and ride to be just off the Lakewood exit as shown in 1.

After running the simulation for a fixed amount of time, we evaluate the average travel time, average carbon emissions, and number of vehicles that reach the sink for the following vehicles:

- cars that intend to travel from source to sink without exiting
- cars that intend to travel from source to sink that are diverted
- buses leaving park and ride structure that drive to sink

We use SUMO's default HBEFA3 emissions model [25]. to obtain emissions values for each vehicle. We use the average speed of all vehicles to evaluate the overall congestion of the scenario. The results are in Table II.

TABLE II: Fixed compliance experimental results

Compliance	Travel Time (s)	CO2 (kg)	Speed ($\frac{m}{s}$)	Count ($\frac{veh}{hr}$)
0%	1063.18	5.33	15.77	1793
10%	1040.60	4.96	16.15	1899
25%	1034.60	4.37	16.53	1992
50%	969.07	3.51	17.48	2188

We can see that increasing the rerouting compliance yields improvements in all categories in terms of average travel time, average CO2 emissions, overall simulation vehicle speed, and throughput. One limitation of this preliminary study is that it cannot distinguish between buses with passengers and buses without passengers, which may inflate the overall throughput and emissions values. Additionally, we are always diverting vehicles in this scenario when in practice we would prefer to divert vehicles in response to downstream congestion.

IV. V2I Communication Proof of Concept

We created a hardware test bed inspired by [26] to demonstrate the V2I communication flow between a smart parking structure, cloud server, RSU, and an OBU for the ICM strategy. To perform wireless DSRC communications, we used the USRP B210 board from Ettus Research, which is a software-defined radio that can receive and transmit data at any frequency between 70 MHz - 6GHz [27]. A diagram of the setup is shown in figure 2.

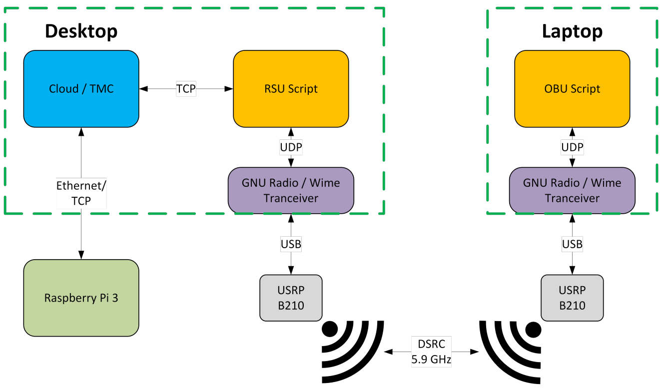

Fig. 2: Block diagram of the hardware test bed

A Raspberry Pi 3 representing a smart parking structure is connected via Ethernet cable to the desktop, which is running a Python script representing a cloud server. These two communicate over TCP. Additionally, in a separate process on the desktop, another Python script is running and represents an RSU. The RSU script communicates to the USRP B210, which will broadcast or listen on the DSRC 5.9GHz band. A picture of this half of the setup is shown in figure 3a. The other half of the setup is a laptop connected to another USRP B210, which can also broadcast or listen on the DSRC 5.9GHz band. A picture of this half of the setup is shown in figure 3b.

To interface with the USRP B210, we used a WiFi tranceiver module created by the open source Wime Project [28], which provides a complete physical layer implementation of 802.11p in GNU Radio. This was modified to send and receive UDP packets from the local machine as shown in figure 4. To send information to the USRP B210 for wireless transmission, the Python script simply writes to a UDP socket at *localhost:52001*. To read information received from the USRP B210, the Python script simply reads from a UDP socket at *localhost:52002*.

A. DSRC Tests

The first test is a forward propagation of data from the Raspberry Pi to the laptop which represents broadcasting park and ride information to a connected vehicle. The Raspberry Pi sends a dictionary of parking space info to the cloud server, which forwards the data to the RSU. The RSU wirelessly broadcasts the message to the awaiting laptop using DSRC. Screenshots demonstrating this data propagation are shown in figures 5a and 5b.

The second test is a backward propagation of data from the laptop to the cloud server which represents collecting state

(a) A RPI 3 connected to the desktop via ethernet, and the B210 plugged into the desktop via USB

(b) The B210 plugged into the laptop via USB

Fig. 3: Hardware setup for testing DSRC communications

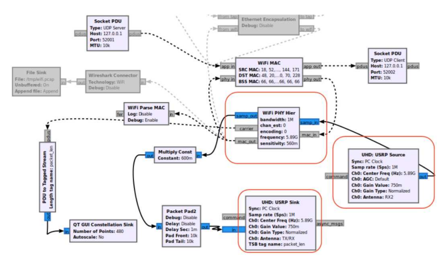

Fig. 4: Hardware blocks for 802.11p in GNU Radio, the same interface from [26]

information from a connected vehicle. The laptop broadcasts some basic vehicle information to the awaiting RSU using DSRC. The RSU then forwards the message to the cloud server.

B. Practical Implementation

This hardware proof of concept clarifies the ICM mechanisms needed to read traffic state and alert drivers to reroute to a PAR structure. In these tests, we send data directly from one script to another, but in practice, standard DSRC message types are needed in both cases. The SAE J2735 standard [29] defines many message types for V2X communications; none are designed specifically for our use case, but a couple message types are flexible enough to be adapted. For an RSU broadcasting an advisory message to reroute to a PAR, one could use the traveler information message (TIM) [29], which is used to broadcast various advisory or road sign info messages. The TIM can be configured to be active on

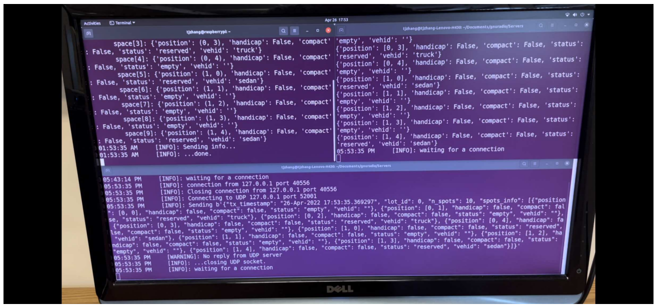

(a) Desktop screenshot showing interface for RPI (top-left), cloud (top-right), and RSU (bottom). Parking info is generated by RPI and propagates to RSU for broadcasting.

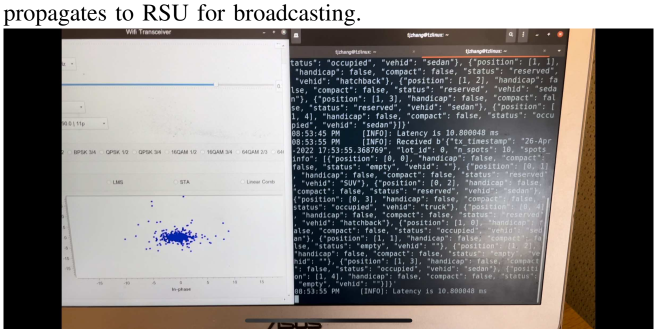

(b) Laptop screenshot showing interface for OBU (right).The OBU receives and displays the parking info from the broadcast.

Fig. 5: Screenshots demonstrating forward flow of information from parking structure to OBU

a minute-by-minute basis and even has limited support for custom strings, which can be useful for informing drivers about the nearest PAR. For collecting state information, an RSU could collect a basic safety messages (BSM) [29] from vehicles nearby and aggregate the data. Part 1 of the BSM frame is mandatory and reports the vehicle's position and velocity. Part 2 of the BSM frame is optional but could be customized with additional information that is of importance to the RSU or ICM strategy, such as emissions information or vehicle type.

The cloud server functions similarly to a Transportation Management Center (TMC) [30] but with the additional role of a park and ride management system. Communicating to the TMC from an RSU or a smart parking structure can be done over LTE; there are examples of RSUs being equipped with bidirectional LTE radios such as one developed by Siemens [31].

V. System Model

In this section, we model the possible travel times experienced by vehicles when traversing the integrated corridor, and we define the evaluation metrics and formulate the optimization problem to effectively manage traffic along the ICM supported by CV technology and PAR facilities.

A. Travel Delay

Assume a vehicle k enters the corridor through a link n with the final destination being the corridor's sink at link N. Given PAR is supported, the time taken by vehicle k to traverse the corridor can be given as follows:

$$T_{k,N} = \begin{cases} \sum_{i=n}^{N-1} \tau_{\pi_i \to \pi_{i+1}}, & \text{if } y_k == 0 \\ \mathcal{T}_{exit} + \mathcal{T}_{PAR}^{total} + \mathcal{T}_{rem}, & \text{if } y_k == 1 \end{cases} \tag{3}$$

where $y_k \in \{0, 1\}$ describes the method by which vehicle k traversed the corridor: y_k=0 indicating the conventional direct approach and y_k=1 implying opting for a PAR option. For the former, the time taken by vehicle k is estimated through the accumulation of times, $\tau_{\pi_i \to \pi_{i+1}}$[1], representing the time taken by the vehicle to traverse from one link i to the next i+1 until the final transition to the sink point, $\tau_{\pi_{N-1} \to \pi_N}$. For the latter case, we breakdown the time experienced by vehicle k commuters opting for a PAR option into three components: (*i*) $\mathcal{T}_{exit}$ is the time taken to reach an exit ramp from the corridor once y_k values has turned to 1, (*ii*) $\mathcal{T}_{PAR}^{total}$ representing the total PAR service time, and (*iii*) $\mathcal{T}_{rem}$ representing the remainder time to be traversed by the returning vehicle until the sink.

B. Evaluation Criteria

Let the set $V_K = \{v_1, v_2, ..., v_K\}$ be the total number of vehicles on the corridor during a time window t, where every $v_k \in V_K$ has entered the corridor from its starting link, π_n, with a final destination at the sink, π_N. Let also $V_{K'} \subset V_K$ be the subset of vehicles that reach the corridor's sink during window t such that $K' \leq K$. As such, we can define the following key metrics to evaluate the overall congestion state along the corridor:

$$\mathbb{FR} = |V_{K'}| \tag{4}$$

$$\mathbb{T} = \frac{\sum_{k=1}^{K'} T_{k,N}}{\mathbb{FR}} \quad \forall \ \ v_k' \in V_{K'} \tag{5}$$

$$\mathbb{CE} = f(V_K, efficiency, SPEED) \tag{6}$$

in which $\mathbb{FR}$ represents the flow rate of the corridor during window t, evaluated as the cardinality of $V_{K'}$; $\mathbb{T}$ represents the average travel time experienced by the vehicles reaching the sink, π_N, evaluated as the sum of individual travel times over the flow rate; $\mathbb{CE}$ represents the carbon emissions exerted by the V_K set of vehicles traversing the corridor during window t. The evaluation of $\mathbb{CE}$ depends on V_K as well as their corresponding fuel efficiency and travel speeds.

C. Problem Formulation

From here, we can formulate our problem in the following manner: within an integrated corridor where RSUs have the capability to broadcast PAR availability to CVs, our aim is to determine the most effective advertising strategy, $X^{*(t)}$, during a specific time t. This strategy should inform upstream traffic about the state of traffic conditions and provide suitable alternatives regarding PAR usage. Thus, we can define the global optimization objective at time t as follows:

$$X^{*(t)} = \max_{X^{(t)}} F(\mathbb{T}, \mathbb{FR}, \mathbb{CE}) \tag{7}$$

[1]We omit the subscript k here for reading convenience

where F represents a global optimization function to be maximized with respect to the evaluation metrics defined in equations 4, 5, and 6. One straightforward implementation of F would be to employ a weighted sum formula with negative weights assigned to the metrics that need to be minimized (i.e., $\mathbb{CE}$ and $\mathbb{T}$). It should be noted that the global optimization objective can be generalized to account for other objectives, such as minimizing the costs of fuel and park and ride charges for commuters.

VI. Experimental Modeling

The basic idea of the experiment is to train the DRL model with the I5 Veins simulation in the loop until the model converges to a policy. Figure 6 shows a high-level overview of the Veins framework. We constructed the Veins simulation based

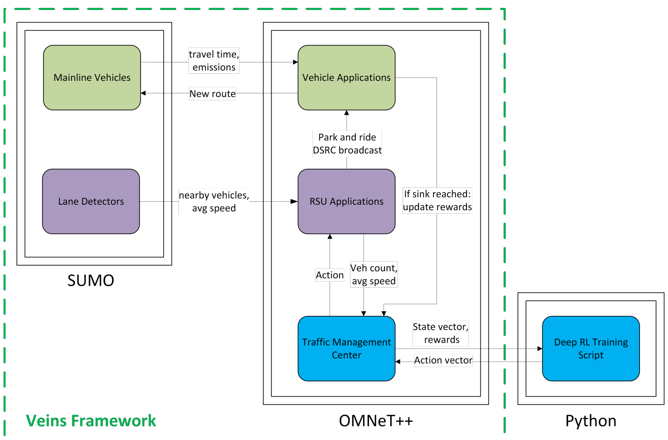

Fig. 6: Flow chart of simulation-in-the-loop training process

on the I5 corridor developed in section 3. For the simulation, we chose to use SUMO's default driving behavior and CO2 emissions model [25]. To reduce computation complexity, We reduced the size of the scenario, setting the sink just after Garfield Ave (see figure 1). SUMO's extensive calculations can only be executed on a single core, creating a bottleneck for the training process. Below are some assumptions made when designing the Veins simulation.

A. Simulation Assumptions

- There is only one PAR structure, the North Lakewood Park and Ride.
- There is only one passenger per vehicle. When vehicles update the accumulated rewards, they will only increase the throughput by one. If HOVs have N passengers onboard, they will increase the throughput by N.
- Commuters rerouted to the PAR structure can take the next HOV leaving the PAR once they arrive. We elaborate on how rerouting delay is computed below.
- Vehicles do not exit the PAR structure during the simulation, the number of available spaces depends solely on the DRL agent.
- HOVs do not encounter additional waiting times at entering ramp meters. We added a dedicated HOV lane at the on ramp nearest to the North Lakewood Park and Ride.
- RSUs are positioned prior to every exit and on ramp along the corridor. See figure 7 to see the full placement.
- DSRC broadcasts have no propagation delay. Simulating broadcast delay would create simulation overhead and does not add much value to the scope of this experiment.
- DSRC range is limited to 75m radius around RSU, which is within the 250-300m effective range of DSRC [32].

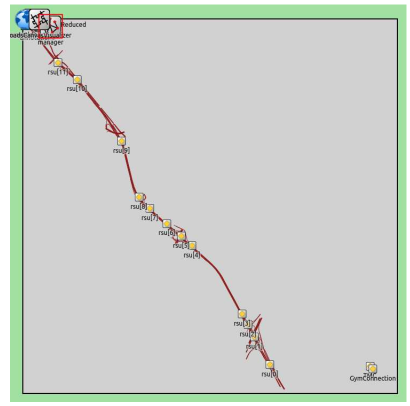

Fig. 7: RSU placement for Veins simulation

B. Simulation Parameters

The following defines a list of Veins simulation parameters.

- S: traffic spawn scale factor
- λ_{HOV}: the frequency of the HOVs leaving the PAR
- t: the time window that constitutes one step in the reinforcement learning algorithm
- α: compliance probability of vehicles with PAR messaging
- γ: probability distribution function capturing additional time for the vehicle to get to and find parking at the PAR
- $f_{redirect}$: the frequency of PAR broadcasts during a time window t
- P: number of parking spaces available at PAR
- r: fixed ramp metering rate for on ramps

We train the DRL agent using the simulation parameters described in table III. The HOV frequency schedule λ_{HOV}

Parameter	Value
S	2.0
λ_{HOV}	$\frac{1}{60sec}$
t	600 s
α	10%
$f_{redirect}$	30 s
P	400
r	900 VPH

TABLE III: Parameter values for Veins simulation

is set to spawn every 60s so passengers will not have to wait excessively compared to the time it takes to travel through the corridor. The amount of time that constitutes a step size and the action space t is set to 600s because it takes approximately 300s for a vehicle to reach the sink and contribute its rewards. The compliance parameter α is deliberately set at a low value of 10% since realistically the majority of drivers would ignore an advisory message to locate the nearest PAR. The ramp metering rate is set to a fixed interval of 900 VPH or one vehicle every four seconds, which is within range of typical ramp meter rates in California [33].

C. Delay

As vehicles exit the simulation, the travel time cost to reach the PAR is computed before its reward is transferred to an HOV or continuing vehicle. The penalty incurred is contingent on the specific exit the vehicle takes. Since we used OpenStreetMap to generate our network as explained in section 3, the simulation distance and speed limits are based on real life values, so we can use Google Maps data to approximate the travel time for exiting vehicles. Using Google Maps, we estimated the travel time to reach the North Lakewood Park and Ride structure from each exit along the corridor. These values are summarized in table IV.

Exit	Time to N Lakewood PAR (s)
I-605	300
Lakewood Blvd	30
Paramount Blvd	180
Slauson Ave	300
Garfield Ave	360
Washington Blvd	540
Atlantic Blvd	720
Atlantic and Triggs	840
Eastern Ave	900
Ditman Ave	1080
Calzona St	1200

TABLE IV: Approximate travel times to reach the North Lakewood Park and Ride from various highway exits, taken from Google Maps

Parameter γ is simply a normal distribution that models the additional time for the vehicle to get to and find parking within the PAR. We take the absolute value of the normal distribution to avoid negative delay values. Thus, the additional travel time penalty for a vehicle taking an exit is represented in equation 8.

$$\gamma = |normal(0, T_{exit}/8)| + (T_{arrival} \bmod \frac{1}{\lambda_{HOV}}) \quad (8)$$

where T_{exit} represents the corresponding exit time from table IV, and $(T_{arrival} \bmod \frac{1}{\lambda_{HOV}})$ represents the additional delay a driver would experience waiting for the next HOV to leave the station. If the vehicle instead finds that there are no more parking spots available, the delay is modeled more harshly. This is shown in equation 9.

$$\gamma = |normal(0, T_{exit}/4)| \quad (9)$$

D. Deep RL Training

From table IV, exits to I-605, Lakewood, Paramount, and Slauson have the smallest rerouting time penalty. Thus, the action vector is restricted to RSU[0], RSU[3], RSU[4], RSU[6], and RSU[7], marked in figure 7. Since each RSU is represented as a bool in the action vector, this limits the action space size to $2^5 = 32$, which helps with exploration.

The DRL agent utilizes two hidden layers for its neural network. Since the state vector involves {average speed, occupancy} for 12 RSUs and the available parking spaces, the flattened input vector is $1 + 12 * 2 = 25$ floating point values. These are fed into a hidden layer of 128 perceptrons using the the ReLU activation function. The next hidden layer is 64 ReLU perceptrons. Finally, the output layer contains 32 perceptrons to approximate the Q values for each possible combination of actions for the 5 RSUs.

In terms of reward function, we programmed the TMC to perform equation 10 when the Veins simulation is running.

$$R = \beta_1 \cdot (\mathbb{FR} - target_{FR}) + \beta_2 \cdot (\mathbb{T} - target_T) + \beta_3 \cdot \mathbb{CE} \quad (10)$$

The reward metric weights are:

- $target_{FR} = 600$. This is the fixed constant that is subtracted from throughput before weighting. Typically at least 600 vehicles can exit the simulation in one simulation step.
- $\beta_1 = \frac{10}{target_{FR}}$. This is the weight for throughput, making it so an additional 60 vehicles will add an extra point of reward.
- $target_T = 400$. This is the fixed constant that is subtracted from travel delay before weighting. Typically it takes at least 400s to reach the end of simulation.
- $\beta_2 = -\frac{1}{target_T}$. This is the weight for travel delay, it is negative so that it will be minimized.
- $\beta_3 = -\frac{4}{10000000}$. This is the weight for the carbon emissions, it is negative so that it will be minimized. The emissions values are in the order of magnitude of 10^6.

We trained the deep learning agent on one fixed seed of the simulation for 1000 episodes to obtain the DRL agent model. Setting the seed to a fixed value makes the traffic behavior (including route probabilities, Poisson process spawning, and driver AI) deterministic within SUMO. Then, in the next section, we evaluate the model's performance on other simulation seeds to test if the model generalizes to different permutations of the defined traffic behavior.

VII. Experiments

To test if the model generalizes to different permutations of the defined traffic behavior, we evaluated the model on several different seeds of the SUMO simulation and plotted the throughput, average travel delay per passenger, and average carbon emissions per passenger for each seed. First, we introduce details of the Veins simulator and in the following sections, we compare these metrics against the same seed simulation without any ICM control agent. In the following figures, we compare the evaluation metrics for our ICM approach against no control for SUMO simulation seeds 87670, 65643, 44435, and 27438. For each seed, we plot the driving duration for mainline vehicles over departure time and the CO2 emissions for mainline vehicles and buses over departure time without the DRL agent. We do the same thing for results with the DRL agent, and in addition, we plot the driving departure versus departure time for rerouted vehicles.

A. VEINS Simulator

In order to simulate V2I communications, we use Veins, an open-source vehicular network simulation framework that is based on two preexisting simulators - OMNet++, an event-based network simulator, and SUMO, a microscopic traffic simulator [34]. Models in OMNeT++ are composed of interconnected components that represent entities or processes in the simulation. It is mainly used for the simulation of communication networks, but it has also been applied to transportation systems, queueing systems, etc [35]. The components of OMNeT++ are reusable and can be combined to create complex models. The modules in OMNeT++ exchange messages, and this helps an individual realize their network of choice. The SUMO simulator is specifically made to model and simulate road traffic in urban regions. It consists of elements such as vehicles, traffic lights, intersections., etc.

Veins is an OMNeT++ project that defines a dynamic network topology from moving SUMO vehicles, models the DSRC communication stack in OMNeT++, and provides an API to control and read values from the underlying SUMO traffic simulation via the Traffic Control Interface (TraCI) [34]. We chose Veins because it is open-source, flexible, and suitable for simulating V2I communications and developing various V2X simulation scenarios.

Finally, we ran one last experiment with large dataset (1000 datapoints in total) and achieved the results shown in Figures **??**, **?? ??**, and **??**. Our DRL agent showed 3 extra vehicles per minute, a reduction of 4 second of average travel time for all vehicles, and a slightly higher average speed along the corridor. However, more buses on the road leads to negligeably higher CO2 emissions (Less than 100g per vehicle).

B. Discussion

We observe that our ICM approach provides marginal improvements to highway throughput, average delay, and average CO2 emissions. Further investigation into the optimal policy learned by the DRL agent shows that the agent primarily redirects traffic to the PAR structure from Interstate 610 and Lakewood Blvd. This makes sense because taking these two exits result in the least amount of travel time to reach the PAR (see table IV). Additionally, rerouting at I-610 may be especially beneficial because as observed in section 3, the merges at the I-610 junction are a heavy source of congestion.

We also find that average delay per passenger does not tell the whole story. Although the travel time for drivers on the corridor is improved, the travel time for drivers using the PAR services is drastically increased in all cases. In other words, by letting a few drivers suffer significantly greater travel delay by taking public transport, the other drivers on the freeway are allowed to get to the destination slightly faster. This is a significant trade-off to keep in mind for this ICM strategy, one which is consistent with the results of Ortega et al.'s study [18], who found that travellers utilizing PAR systems experience significantly increased travel time.

Our DRL agent slightly increase the CO2 emissions however, this is likely due to the fact that HOVs are not utilized in the uncontrolled scenario but still contribute to the overall

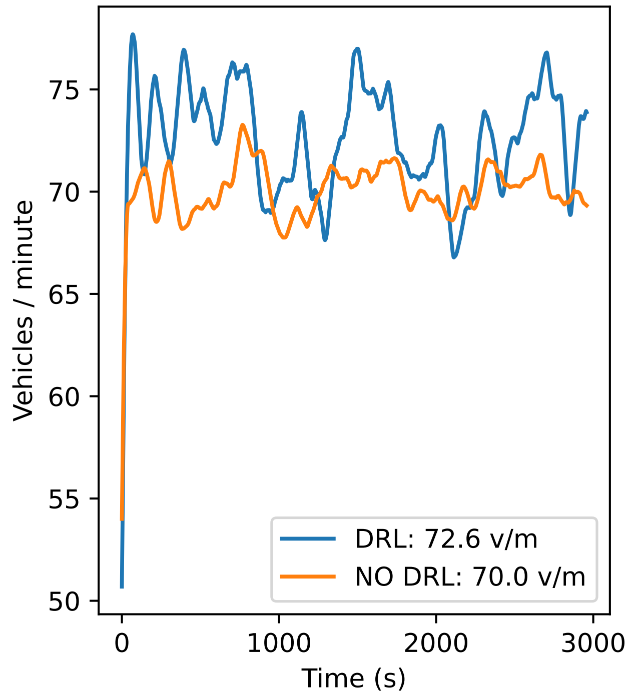

Fig. 8: Flowrate vs Time

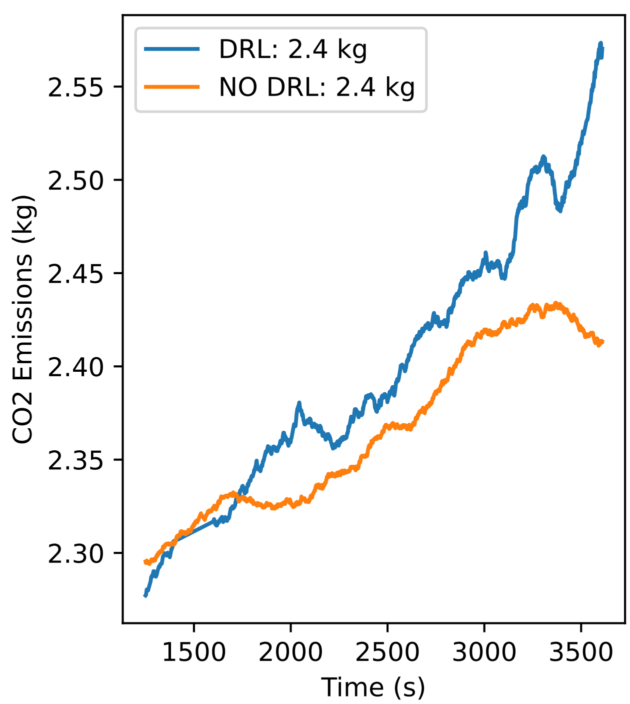

Fig. 9: CO2 Emissions vs Time

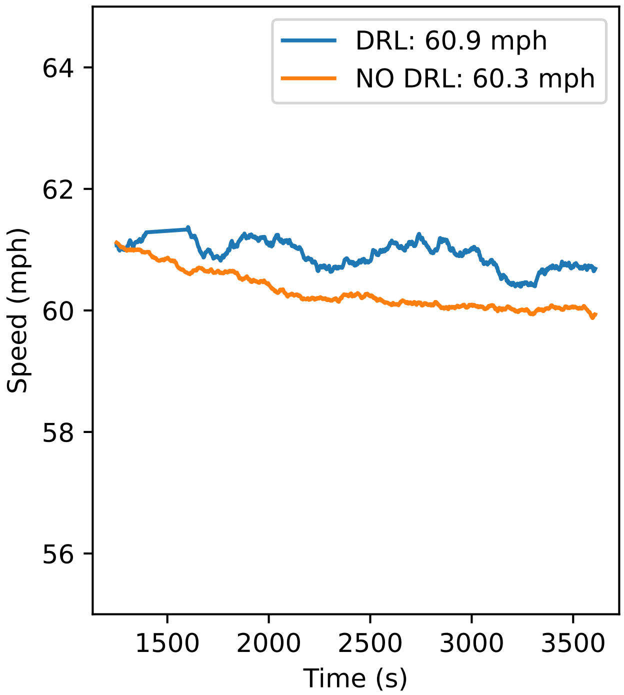

Fig. 10: Speed vs Time

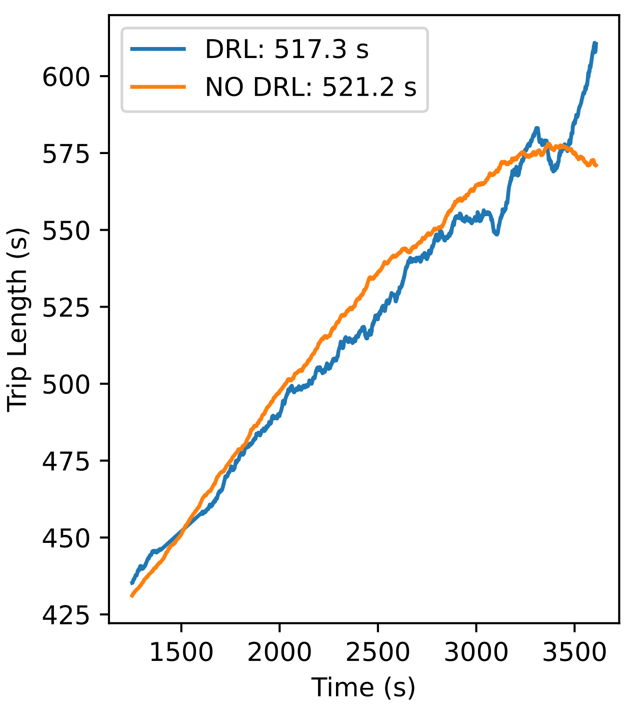

Fig. 11: Trip Length vs Time

emissions of the corridor, wasting resources. When passengers ride on an HOV, the CO2 cost for the vehicle per passenger decreases at a rate of $\frac{1}{N}$, rapidly increasing the efficiency of the HOV. Modifying some simulation parameters could result in even greater HOV emission savings. For instance, a less frequent HOV schedule $\lambda_{HOV,j}$ would reflect more realistic public transportation schedules and allow more passengers to accumulate at the PAR before the next HOV leaves, but this would further sacrifice the travelers' delay during the extended waiting period. Another possibility is to set a higher reroute compliance probability α to reroute more vehicles to the PAR, but too high a value would be unrealistic; not many drivers are willing to reroute to a PAR. However, a higher α could be made a realistic assumption if drivers are piloting autonomous vehicles instead of just connected vehicles.

VIII. Future Work

In the Discussion section, we recommend some parameter adjustments, such as decreasing the HOV schedule frequency and increasing driver compliance with PAR messages, to see if there are scenarios where rerouting vehicles to the PAR structure can create even more emissions savings. These changes could be translated into a study involving autonomous vehicles, where the system could have more control over the behavior of autonomous passenger vehicles and the timing of autonomous HOVs.

More studies are needed to see how this approach can be scaled up. The difficulty in this will be in observing and assigning rewards to the ICM actions; with our current experimental setup, a longer freeway means more time will pass before a passenger reaches the sink in an HOV, i.e. more time will pass before rewards reflect the new action. This increases simulation complexity and overall training time. One alternative could be to deploy multiple DRL agents for small sections of the freeway and coordinate them into a larger system.

Additionally, as with any system that aggregates data to make decisions, there are security concerns. An automotive security survey conducted in 2019 [36] explains that V2X communications opens up multiple new attack surfaces for vehicles in addition to the preexisting vulnerabilities in automotive electronic components. In particular, the authors find data spoofing to be a common attack method on V2I-based systems such as our ICM approach, resulting in increased traffic congestion. A 2022 study [37] develops an attack modeling methodology for a V2X Advisory Speed Limit traffic control scenario and establishes various metrics to assess the impact of an attack. A future research direction could be to develop a similar attack methodology and evaluation metrics for our V2I-based ICM approach and to test its resilience.

One could also adopt a more general security framework to analyze our ICM approach. In a previous study [38], the authors present a security analysis framework for cyber-physical systems (CPS). By modeling cyber domain information, such as device firmware and application data, and physical domain information, such as RF signals and other side channels, as information flows, the authors show that applying data-driven algorithms can improve understanding of the cyber-physical domain relationships and reveal new vulnerabilities in the system. Moreover, our ICM approach can be classified as a networked control system (NCS) that can be modeled as generalized mathematical formula as demonstrated in [39]. Once the vulnerabilities and attacks of our ICM strategy are understood, they can be modeled and fed as input to the NCS model to study the system response to attacks over time.

Any proposed security solutions should focus on extensibility [40], i.e. solutions that are easily adapted to new use cases to support the still-developing automotive technology scene. Several previous works describe extensible security solutions for V2X communications. For instance, the authors in [41] and [42] propose novel methods for physical layer key generation that result in faster key generation time and reduced code size respectively; reduced computational resources means these cryptography methods can be implemented in more devices. In another study [43], the authors propose a blockchain-based architecture to validate a connected vehicle's location in V2I contexts, preventing position spoofing. These solutions can be incorporated into our ICM strategy.

IX. Conclusion

This work proposes a novel ICM strategy that redirects vehicles to underutilized park and ride structures to maximize freeway throughput and minimize CO2 emissions and travel time. This approach leverages the V2I capabilities of RSUs and OBUs to observe the state of connected vehicles on the freeway and to broadcast advisory messages to drivers to redirect them to the nearest park and ride structure. A centralized cloud server hosted at a Transporation Management Center communicates with the RSUs and uses deep reinforcement learning to process the observed congestion state of the corridor and choose where to broadcast PAR advisory messages.

We created a realistic corridor simulation based on Interstate 5 in the Los Angeles area to test the ICM strategy. The deep reinforcement learning agent converges to a strategy that redirects vehicles to the I-605 and Lakewood Blvd junctions, which can achieve marginal improvements in throughput, average travel time, and average emissions at the cost of significant travel delay for the few drivers taking an HOV. Specifically, we observe up to 3.99% increase in throughput, 4.67% reduction in freeway travel time, and 3.09% savings in CO2 emission savings, but with the cost of up to 52.56% additional delay for diverted drivers.

Acknowledgements

This study was facilitated by funding obtained by the University of California Institute of Transportation Studies from the State of California via the Public Transportation Account and the Road Repair and Accountability Act of 2017 (Senate Bill 1). The contents of this report reflect the views of the author(s), who hold accountability for the facts and accuracy of the information provided. This document is disseminated under the sponsorship of the State of California in the interest of information exchange and does not necessarily reflect the official views or policies of the State of California.

References

[1] D. Schrank, L. Albert, B. Eisele *et al.*, "Urban mobility report 2021," 2021.

[2] INRIX, "2021 global traffic scorecard," 2021. [Online]. Available: https://inrix.com/scorecard/

[3] F. H. Administration, "Integrated corridor management (icm) program: Major achievements, key findings, and outlook," https://ops.fhwa.dot.gov/publications/fhwahop19016/chapter3.htm, 2020.

[4] ——, "Integrated corridor management, transit, and mobility on demand," https://ops.fhwa.dot.gov/publications/fhwahop16036/ch1.htm, 2020.

[5] Caltrans, "Park & ride," 2023. [Online]. Available: https://dot.ca.gov/programs/traffic-operations/park-ride

[6] Transportation Research Board and National Academies of Sciences, Engineering, and Medicine, *Transit Supportive Parking Policies and Programs*, L. Jacobson and R. R. Weinberger, Eds. Washington DC: The National Academies Press, 2016. [Online]. Available: https://nap.nationalacademies.org/catalog/23493/transit-supportive-parking-policies-and-programs

[7] A. Schmitt, "Park & rides lose money and waste land - but agencies keep building them," Jul 2016. [Online]. Available: https://usa.streetsblog.org/2016/07/05/park-rides-lose-money-and-waste-land-but-agencies-keep-building-them/

[8] "Dedicated short range communications (dsrc) service," 2022. [Online]. Available: https://www.fcc.gov/wireless/bureau-divisions/mobility-division/dedicated-short-range-communications-dsrc-service

[9] "Roadside and on-board unit in-kind grant," 2022. [Online]. Available: https://www.codot.gov/programs/innovativemobility/assets/oim-grants/rsu_obu_grant_rules-selection_criteria.pdf

[10] J. Hu, P. Bhowmick, F. Arvin, A. Lanzon, and B. Lennox, "Cooperative control of heterogeneous connected vehicle platoons: An adaptive leader-following approach," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 977–984, April 2020.

[11] J. Shelton, J. Wagner, S. Samant, G. Goodin, T. Lomax, and E. Seymour, "Impacts of connected vehicles in a complex, congested urban freeway setting using multi-resolution modeling methods," *International Journal of Transportation Science and Technology*, vol. 8, no. 1, pp. 25–34, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2046043018300352

[12] W. Jin and M. Zhang, "Evaluation of On-ramp Control Algorithms," Institute of Transportation Studies, UC Berkeley, Institute of Transportation Studies, Research Reports, Working Papers, Proceedings qt1gz7w0wm, Apr. 2001. [Online]. Available: https://ideas.repec.org/p/cdl/itsrrp/qt1gz7w0wm.html

[13] A. Fares and W. Gomaa, "Freeway ramp-metering control based on reinforcement learning," *IEEE International Conference on Control and Automation, ICCA*, 06 2014.

[14] H. Hashemi and K. F. Abdelghany, "Real-time traffic network state estimation and prediction with

decision support capabilities: Application to integrated corridor management," *Transportation Research Part C: Emerging Technologies*, vol. 73, pp. 128–146, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0968090X16302054

[15] H. Hashemi and K. Abdelghany, "End-to-end deep learning methodology for real-time traffic network management," *Computer-Aided Civil and Infrastructure Engineering*, vol. 33, no. 10, pp. 849–863, 2018. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/mice.12376

[16] Y. Han, M. Wang, L. Li, C. Roncoli, J. Gao, and P. Liu, "A physics-informed reinforcement learning-based strategy for local and coordinated ramp metering," *Transportation Research Part C: Emerging Technologies*, vol. 137, p. 103584, 2022.

[17] X. C. Liu, G. Zhang, C. Kwan, Y. Wang, and B. K. Kemper, "Simulation-based, scenario-driven integrated corridor management strategy analysis," *Transportation Research Record*, vol. 2396, no. 1, pp. 38–44, 2013. [Online]. Available: https://doi.org/10.3141/2396-05

[18] J. Ortega, J. Hamadneh, D. Esztergár-Kiss, and J. Tóth, "Simulation of the daily activity plans of travelers using the park-and-ride system and autonomous vehicles: Work and shopping trip purposes," *Applied Sciences*, vol. 10, no. 8, 2020. [Online]. Available: https://www.mdpi.com/2076-3417/10/8/2912

[19] J. Hamadneh and D. Esztergár-Kiss, "Travel behavior of car travelers with the presence of park-and-ride facilities and autonomous vehicles," *Periodica Polytechnica Transportation Engineering*, vol. 50, no. 1, pp. 101–110, 2022.

[20] X. Liu, N. Masoud, Q. Zhu, and A. Khojandi, "A markov decision process framework to incorporate network-level data in motion planning for connected and automated vehicles," *Transportation Research Part C: Emerging Technologies*, vol. 136, p. 103550, 2022.

[21] "L.A. County's I-5 named 'most congested freeway' in California," 2014. [Online]. Available: https://www.dailynews.com/2014/02/13/la-countys-i-5-named-most-congested-freeway-in-california/

[22] S. Carpenter, "Interstate 5 from Euclid Avenue to 605 is busiest corridor in the U.S." *Spectrum News 1*, 2021. [Online]. Available: https://spectrumnews1.com/ca/la-west/traffic/2021/12/07/interstate-5-from-euclid-avenue-to-interstate-605-is-busiest-corridor-in-the-u-s-

[23] "Los Angeles Has Top Two Worst Traffic Corridors in the Nation, Study Says ," 2020. [Online]. Available: https://www.nbclosangeles.com/news/local/los-angeles-has-top-two-worst-traffic-corridors-in-the-nation-study-says/2325974/

[24] "Traffic Census Program." [Online]. Available: https://dot.ca.gov/programs/traffic-operations/census

[25] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using sumo," in *The 21st IEEE International Conference on Intelligent Transportation Systems*. IEEE, 2018. [Online]. Available: https://elib.dlr.de/124092/

[26] M. A. Al Faruque, M. Odema, and L. Chen, "Software and hardware systems for autonomous smart parking accommodating both traditional and autonomous vehicles," Pacific Southwest Region University Transportation Center (UTC), Tech. Rep. PSR-19-30, Apr. 2021. [Online]. Available: https://rosap.ntl.bts.gov/view/dot/56844

[27] E. Research, "Usrp hardware driver and usrp manual," https://files.ettus.com/manual/page_usrp_b200.html, 2023.

[28] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, "Performance Assessment of IEEE 802.11p with an Open Source SDR-based Prototype," *IEEE Transactions on Mobile Computing*, vol. 17, no. 5, pp. 1162–1175, May 2018.

[29] V. C. T. Committee, *V2X Communications Message Set Dictionary*, jul 2020. [Online]. Available: https://doi.org/10.4271/J2735_202007

[30] F. H. Administration, "Transportation management centers," https://ops.fhwa.dot.gov/freewaymgmt/trans_mgmnt.htm, 2020.

[31] Siemens, "Connected vehicle roadside unit (rsu)," https://www.mobotrex.com/product/siemens-connected-vehicle-roadside-unit/, 2018.

[32] F. Arena, G. Pau, and A. Severino, "A review on ieee 802.11p for intelligent transportation systems," *Journal of Sensor and Actuator Networks*, vol. 9, no. 2, 2020. [Online]. Available: https://www.mdpi.com/2224-2708/9/2/22

[33] D. of Traffic Operations, "Ramp metering design manual," https://dot.ca.gov/programs/traffic-operations/ramp-metering, California Department of Transportation, Tech. Rep., 2022.

[34] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing (TMC)*, vol. 10, no. 1, pp. 3–15, January 2011.

[35] A. Varga, *OMNeT++*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 35–59. [Online]. Available: https://doi.org/10.1007/978-3-642-12331-3_3

[36] A. Lopez, A. V. Malawade, M. A. Al Faruque, S. Boddupalli, and S. Ray, "Security of emergent automotive systems: A tutorial introduction and perspectives on practice," *IEEE Design & Test*, vol. 36, no. 6, pp. 10–38, 2019.

[37] A. B. Lopez, W.-L. Jin, and M. A. A. Faruque, "Attack modeling methodology and taxonomy for intelligent transportation systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 13 255–13 264, 2022.

[38] S. R. Chhetri, J. Wan, and M. A. Al Faruque, "Cross-domain security of cyber-physical systems," in *2017 22nd Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2017, pp. 200–205.

[39] A. Sargolzaei, A. Abbaspour, M. A. Al Faruque, A. Salah Eddin, and K. Yen, *Security Challenges*

of Networked Control Systems. Cham: Springer International Publishing, 2018, pp. 77–95. [Online]. Available: https://doi.org/10.1007/978-3-319-74412-4_6

[40] S. Ray, W. Chen, J. Bhadra, and M. A. Al Faruque, "Extensibility in automotive security: Current practice and challenges: Invited," in *Proceedings of the 54th Annual Design Automation Conference 2017*, ser. DAC '17. New York, NY, USA: Association for Computing Machinery, 2017. [Online]. Available: https://doi.org/10.1145/3061639.3072952

[41] S. Ribouh, K. Phan, A. V. Malawade, Y. Elhillali, A. Rivenq, and M. A. A. Faruque, "Channel state information-based cryptographic key generation for intelligent transportation systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 12, pp. 7496–7507, 2021.

[42] J. Wan, A. Lopez, and M. A. A. Faruque, "Physical layer key generation: Securing wireless communication in automotive cyber-physical systems," *ACM Trans. Cyber-Phys. Syst.*, vol. 3, no. 2, oct 2018. [Online]. Available: https://doi.org/10.1145/3140257

[43] A. Didouh, A. B. Lopez, Y. E. Hillali, A. Rivenq, and M. A. A. Faruque, "Eve, you shall not get access! a cyber-physical blockchain architecture for electronic toll collection security," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, 2020, pp. 1–7.