

Safe and Efficient CAV Lane Changing using Decentralised Safety Shields

Bharathkumar Hegde Mélanie Bourroche

School of Computer Science and Statistics, Trinity College Dublin, Ireland
{hegdeb, Melanie.Bourroche}@tcd.ie

Abstract—Lane changing is a complex decision-making problem for Connected and Autonomous Vehicles (CAVs) as it requires balancing traffic efficiency with safety. Although traffic efficiency can be improved by using vehicular communication for training lane change controllers using Multi-Agent Reinforcement Learning (MARL), ensuring safety is difficult. To address this issue, we propose a decentralised Hybrid Safety Shield (HSS) that combines optimisation and a rule-based approach to guarantee safety. Our method applies control barrier functions to constrain longitudinal and lateral control inputs of a CAV to ensure safe manoeuvres. Additionally, we present an architecture to integrate HSS with MARL, called MARL-HSS, to improve traffic efficiency while ensuring safety. We evaluate MARL-HSS using a gym-like environment that simulates an on-ramp merging scenario with two levels of traffic densities, such as light and moderate densities. The results show that HSS provides a safety guarantee by strictly enforcing a dynamic safety constraint defined on a time headway, even in moderate traffic density that offers challenging lane change scenarios. Moreover, the proposed method learns stable policies compared to the baseline, a state-of-the-art MARL lane change controller without a safety shield. Further policy evaluation shows that our method achieves a balance between safety and traffic efficiency with zero crashes and comparable average speeds in light and moderate traffic densities.

Index Terms—Connected Autonomous Vehicles (CAVs), Multi-Agent Reinforcement Learning (MARL), Control Barrier Functions (CBFs), Safe control, Lane change

I. INTRODUCTION

Autonomous Vehicles (AVs) were expected to be commercially available by 2020, but recent reports suggest that wider adoption of AVs can only be expected after 2030 or beyond due to societal, regulatory, and technical challenges [1]. Complex technical problems, such as localisation, mapping, perception, route planning, and motion control, are yet to be solved to enable commercial AV deployments [2]. Among them, safe and efficient lane change in dynamic traffic conditions is one of the difficult problems because AV controllers must analyse the environment and actuate a control decision in a very short time period, almost instantaneously. Moreover, improper lane change can lead to a collision that could cause significant damage and loss of lives. Safe and strategic lane changes can improve traffic efficiency in scenarios that require mandatory lane changes, such as highway merging [3].

Using vehicular communication (V2X), Connected Autonomous Vehicles (CAVs) can collaborate to analyse the environment and make efficient control decisions. In particular, CAVs can make efficient lane change decisions using

Multi-Agent Reinforcement Learning (MARL) [4], [5]. Many forms of MARL have been applied for designing lane change controllers including Deep Q-Networks (DQNs) [6], [7], [8], and Actor-Critic Networks (ACN) [9], [10], [11]. Among them, MARL-CAV is an open-source state-of-the-art MARL lane change controller designed for CAVs [11]. MARL-CAV formulates a Partially Observable Markov Decision Process (POMDP) to train lane change controllers in a highway merging section. This approach uses a predication-based priority assignment to avoid collisions, but the safety of the lane change is not guaranteed. Moreover, this priority-based safety layer requires a centralised server to monitor and communicate with CAVs, which hinders its scalability.

In our previous work, we found that Control Barrier Functions (CBFs) are suitable to ensure the safety of CAV lane change controllers [12]. CBFs have also been applied to ensure the safe operation of single-agent RL-based AV controllers [13]. This approach demonstrates the longitudinal safety of a vehicle using simple kinematic equations. Furthermore, Han et al. demonstrate that CBFs can be used to define longitudinal and lateral safety constraints for MARL-based CAV controllers [14]. Their CBF formulation, however, can be computationally expensive as reference trajectories are calculated for all possible behavioural control actions. To improve computational efficiency while constraining both longitudinal and lateral control of an AV, the use of Higher-Order CBFs (HOCBFs) has been proposed [15]. Notably, the lateral constraints in this formulation assume that the rate of heading change is directly controlled. While this assumption is appropriate for simple robotic systems or simulation setups, AVs are typically controlled by adjusting the steering angle rather than the heading directly. Moreover, existing CBF formulations define safety constraints to maintain a constant distance from obstacles. However, using a dynamic safe distance based on time headway can improve traffic efficiency [16].

Overall, the existing work does not fully address the challenges of making safe and efficient lane changes. While MARL lane change controllers can be trained to perform efficient lane changes, they do not guarantee safety. Although CBFs can be integrated with MARL to ensure safety, existing approaches do not fully address safe lateral control and do not support dynamic safety distances. In our previous work [17], we proposed a theoretical framework, CBF-CAV, that integrates CBF constraints with MARL to ensure safe exe-

cution of the behavioural lane change decisions, but it was not empirically evaluated. This paper significantly extends the CBF-CAV framework with the following contributions:

- We propose a decentralised Hybrid Safety Shield (HSS) to constrain CAV controls using the CBFs while respecting their control limits. This shield is designed to effectively ensure the safety of both longitudinal and lateral control inputs based on a dynamic safety constraint.
- We define a MARL-HSS architecture to integrate MARL with the HSS for efficiently training CAV lane change controllers with a safety guarantee from HSS. With MARL-HSS, CAV lane change controllers can be trained to make efficient decisions while ensuring safety.
- The MARL-HSS is comprehensively evaluated in a simulation of a highway on-ramp scenario with different traffic densities. The results show that MARL-HSS outperforms the state-of-the-art MARL lane change controller by effectively balancing between safety and efficiency.

The proposed method, MARL-HSS, considers pure CAV traffic, with vehicles capable of sharing their state without any communication delays or packet loss. This assumption avoids the risk of unknown behaviour from other vehicles, which is not handled in the current safety shield.

II. BACKGROUND

This section introduces the background knowledge on AV control hierarchies, MARL, vehicle model, and CBFs to provide a context for the proposed method.

A. Control hierarchies

The control decisions of AVs can be separated into four hierarchical control levels such as route planning, behavioural, motion planning, and vehicular control [18]. The route planning layer first identifies a feasible route to the destination provided by the user using road network information. To move along the defined route, a behavioural layer makes discrete decisions to change lanes or follow the current lane. The motion planning layer then identifies the control references to execute the desired driving manoeuvre to implement each behavioural decision. These driving manoeuvres are then actuated by the vehicular controller. Among them, control decisions from high-level behavioural and low-level motion planning layers influence the safety and efficiency of lane changes. Therefore, this article mainly focuses on designing these layers in the AV control hierarchy. Moreover, since connectivity among AVs enables collaborative control decisions to improve traffic efficiency, this article considers CAVs for their potential benefits [19].

B. MARL for on-ramp merging

CAVs can use MARL to make efficient lane change decisions at a behavioural level based on its state and the states of surrounding vehicles to improve the traffic flow. MARL-CAV is a state-of-the-art decentralised MARL formulation for on-ramp merging, where each vehicle observes only a part of the environment [11]. Therefore, a POMDP can be formulated

for a multi-agent system of $\mathcal{N}$ vehicles described by a tuple $(\{\mathcal{S}_i, \mathcal{A}_i, \mathcal{R}_i\}_{i \in \mathcal{N}}, \mathcal{T})$.

The *state space*, $\mathcal{S}_i$, of a vehicle i consists of following state variables,

- `is_present` : a binary variable representing the presence of a vehicle within the considered scenario.
- x : The longitudinal position of the vehicle.
- y : The lateral position of the vehicle.
- v_x : The longitudinal velocity of the vehicle.
- v_y : The lateral velocity of the vehicle.
- ψ : The vehicle heading with respect to the road.

The state space also includes the states of n surrounding vehicles within the perception range of a CAV. While the state variables of the ego vehicle are captured with respect to a global coordinate system, the observed vehicle's state variables are relative to the ego vehicle. A CAV can measure its own state from onboard sensors such as LIDAR, RADAR, Camera, GPS, and IMU. These sensors are assumed to provide accurate state estimation. The V2X communication interface can be used to capture the states of the surrounding vehicles. The overall multi-agent state space is defined as a Cartesian product of the individual states $\mathcal{S} = \mathcal{S}_0 \times \mathcal{S}_1 \times \mathcal{S}_2 \times \dots \times \mathcal{S}_{\mathcal{N}}$.

The *action space*, $\mathcal{A}_i$, consists of five discrete variables representing a specific lane change behaviour such as $\{\text{right, left, follow lane, speed up, slow down}\}$. The behavioural layer chooses one of these high-level actions. The overall multi-agent action space is defined as joint actions from the individual AVs, $\mathcal{A} = \mathcal{A}_0 \times \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_{\mathcal{N}}$.

The *reward function*, $\mathcal{R}_i$, rewards an individual CAV for achieving individual and collaborative goals. The reward for a CAV i can be defined as

$$r_i = w_c r_c + w_s r_s + w_h r_h + w_m r_m \quad (1)$$

where a CAV is rewarded for avoiding collision with r_c , maintaining desirable speed with r_s , maintaining desirable headway with r_h , and minimising the merging wait time with r_m . Each of these rewards is multiplied with a weight w_* that can be tuned to prioritise the different objectives. The individual rewards from the n observed vehicles within the perception range are combined to consider collaborative goals such as traffic efficiency. Therefore, the reward function, $\mathcal{R}_i$, is defined as

$$R_i = \frac{1}{n} \sum_{j=0}^n r_j \quad (2)$$

The transition probability, $\mathcal{T}$, characterises the dynamics of the CAV multi-agent system. To navigate CAVs based on control inputs, a vehicle model (explained in section II-C) is used. This MARL formulation is integrated with a safety shield proposed in Section III.

C. Vehicle model

It is difficult to model AV dynamics as they are affected by a variety of factors ranging from vehicle control to weather conditions. For simplicity, the kinematic bicycle model is




Fig. 1. Kinematic bicycle model

often used to model AV motion [20], and we demonstrate the proposed safety shield using this model.

The kinematic bicycle model considers the two front wheels as one wheel and the same for the back wheels, as illustrated in Fig. 1. The following equations define the vehicle model based on the state variables defined previously (II-B),

$$\begin{aligned} \dot{x} &= v_x, & \dot{y} &= v_y, & \dot{v}_x &= a \cos(\psi + \beta), \\ \dot{v}_y &= a \sin(\psi + \beta), & \dot{\psi} &= \frac{2v}{V_l} \sin \beta, \end{aligned} \quad (3)$$

where β is a slip angle at the centre of gravity evaluated based on the steering angle input δ :

$$\beta = \tan^{-1} \left(\frac{1}{2} \tan \delta \right) \quad (4)$$

This vehicle model is used to integrate MARL with the CBF constraints formulated within the proposed HSS in Section III-B.

D. Control Barrier Functions for motion planning

A CBF can be used to constrain a non-linear control system such as the motion planning layer in a vehicle control to ensure safety [21]. Consider a discrete time non-linear control system defined by the following transition dynamics

$$s_{t+1} = f(s_t) + g(s_t)u_t \quad (5)$$

where the next state s_{t+1} is defined based on unactuated dynamics $f : S \rightarrow S$, and actuated dynamics $g : S \rightarrow \mathbb{R}^{n,m}$, n and m are the number of variables in the state space S and action space U respectively, $s_t \in S$ is the system state, and $u_t \in U$ is the control action at time step t . The f and g are defined based on known system dynamics and they are locally Lipschitz continuous, in other words, continuous functions limited by a maximum rate of change.

Consider a safe set C defined as the super-level set of the continuously differentiable function $h : S \rightarrow \mathbb{R}$

$$C : \{s_t \in S : h(s_t) \geq 0\} \quad (6)$$

To ensure the safety of the control system (5), the safe set C must be forward invariant, which means that safe actions can be defined for each state $s_t \in C$ such that the system continues to stay in C . The safe set C is forward invariant if

the function h is a *Control Barrier Function* (CBF) such that there exists $\eta \in [0, 1]$ for all $s_t \in C$ satisfying the following equation

$$\sup_{u_t \in U} [h(f(s_t) + g(s_t)u_t) + (\eta - 1)h(s_t)] \geq 0 \quad (7)$$

where η is the coefficient of zero order control that defines the magnitude at which the system is pushed within the safe set C [13]. Using smaller values of η can enforce the constraints strictly, whereas higher values can relax the constraints. Therefore, η represents how strongly the barrier function pushes the states inwards within C . The existence of a CBF implies that for all $s_t \in C$, there exist u_t such that C is forward invariant [21]. Therefore, the goal is to find a minimal safe action u_t that satisfies (7) to ensure the safety of a control system (5).

To define a CBF h , let us consider the affine barrier function of the form

$$h(s_t) = p^T s_t + q \quad (8)$$

where $p \in \mathbb{R}^n$ and $q \in \mathbb{R}$ are the parameters used to define a safety constraint h on the state s_t . Combining the affine barrier function with the condition (7), the following constraint can be defined for the control action u_t ,

$$-p^T g(s_t)u_t \leq p^T f(s_t) + p^T(\eta - 1)s_t + \eta q \quad (9)$$

To consider multiple safety constraints defined using CBFs, C can be considered as the intersecting half spaces defined by k affine barrier functions [15]. The affine constraint on u_t can be defined by stacking all the constraints.

$$\begin{aligned} Au_t &\leq b, \\ \text{where, } A &= [-p_1^T g(s_t), \dots, -p_k^T g(s_t)] \\ b &= [b_1, \dots, b_k], \text{ with } b_i = p_i^T f(s_t) + \\ &\quad p_i^T(\eta - 1)s_t + \eta q_i \end{aligned} \quad (10)$$

This constraint can be used to reformulate the CBF given by Eq. (7) into the following optimisation problem with a slack variable ϵ in the safety condition and a large constant, K_ϵ that penalizes safety violations.

$$\begin{aligned} u_t &= \arg \min_{u_t} \|u_t\|_2 + K_\epsilon \epsilon \\ \text{s.t. } Au_t &\leq b, \end{aligned} \quad (11)$$

This optimisation problem can be efficiently solved in each time step using Quadratic Program (QP) [22].

If the CBF h is defined for state variables such as x and y , control input u_t cannot be effectively constrained using the above optimisation problem (11) because the constraints are effective only when A is non-zero. The parameter A , defined in Eq. (10), consists of $-p_i^T g(s_t)$, where $g(s_t)$ can be zero if a state variable is not affected by the control input u_t directly. Since acceleration control a_t changes the velocity v of the vehicle in the immediate next step, $g(s_t)$ will be a non-zero value for the CBF constraints defined for velocities v_x and v_y . However, $g(s_t)$ will be zero for the constraints defined on the position of a vehicle, such as x and y . Since the proposed Hybrid Safety Shield (III) defines constraints on the vehicle's position, it uses the velocity v as a control input.




Fig. 2. Topology to identify surrounding vehicles

III. PROPOSED APPROACH

In this section, a Hybrid Safety Shield formulation for a multi-agent system of CAVs is defined using CBF constraints. First, the HSS formulation is presented to constrain control inputs of the motion planning layer to ensure safety. Then, MARL-HSS architecture is presented to integrate HSS with a MARL behavioural lane change controller. The MARL-HSS aims to maximise traffic efficiency while ensuring safety.

A. Hybrid Safety Shield

We propose a Hybrid Safety Shield that constrains longitudinal and lateral control inputs while considering their physical limits using CBFs to ensure safety. In the HSS, an optimisation problem derived from Eq. (11) is solved to ensure safe longitudinal motion and a rule-based method is used to ensure safe lateral motion. The optimisation approach uses CBF constraints defined based on the longitudinal time headway with respect to a leading vehicle. The rule-based approach depends on longitudinal time headways with respect to the vehicles in the target lane. Moreover, the control input's physical limits are integrated into the optimisation problem.

Identifying the surrounding vehicles is important to make safe manoeuvres. Therefore, we define a topology to identify the surrounding vehicles in a scenario with two lanes. This topology identifies three observed CAVs C_{o*} , namely the *leading* vehicle C_{ol} , the *target leading* vehicle C_{otl} , and the *target rear* vehicle C_{otr} . The vehicle immediately in front of the ego vehicle in the same lane is considered to be the leading vehicle. While changing lanes, a vehicle moving immediately in front of the ego vehicle in the target lane is considered the target leading vehicle. A rear vehicle in the target lane is considered to be the target rear vehicle. In a merging section with two lanes, the topology of the surrounding vehicles is shown in Fig. 2. The longitudinal distances from the surrounding vehicles C_{ol} , C_{otl} , and C_{otr} are defined as Δx_{ol} , Δx_{otl} , and Δx_{otr} respectively. The surrounding vehicle states are used in the safety constraints defined for longitudinal and lateral vehicle controls.

1) *Longitudinal control*: Let us define a control system (5) for the longitudinal motion of a vehicle as follows

$$x_{t+1} = x_t + v \cos(\psi + \beta) \quad (12)$$

where x_t is the position along the longitudinal direction at time step t , and ψ and β are the heading angle and slip angle defined in the vehicle model (3). This Eq. (12) can be extended to a multi-agent system of CAVs in a matrix form as follows,

$$\begin{bmatrix} x_e \\ x_o \end{bmatrix} = \underbrace{\begin{bmatrix} x_e \\ x_o \end{bmatrix}}_{f(x)} + \underbrace{\begin{bmatrix} \cos(\psi + \beta) \\ 1 \end{bmatrix}}_{g(x)} \underbrace{\begin{bmatrix} v_e & v_o \end{bmatrix}}_{u_t} \quad (13)$$

where x_e and v_e are the longitudinal position and velocity of the ego vehicle. Similarly, x_o and v_o are the longitudinal position and velocity of an observed vehicle. The velocity of an observed vehicle is considered in the control system to account for its state changes when it makes the worst-case control decision. For simplicity, the time step subscript t is omitted in the above equation.

Since the control input, v_e , can be possibly unsafe, our aim is to evaluate a minimum correction value v^{cbf} to comply with a CBF safety constraint. Therefore, v^{safe} can be defined as

$$v^{safe} = \underbrace{\begin{bmatrix} v_e & v_o \end{bmatrix}}_v + \underbrace{\begin{bmatrix} v_e^{cbf} & 0 \end{bmatrix}}_{v^{cbf}} \quad (14)$$

In the v^{cbf} , the safe correction value for the observed vehicle is set to 0 as its speed cannot be controlled by the ego vehicle.

The QP optimisation problem (11) can be updated using Eq. (14) to evaluate a minimum value for v^{cbf} as shown below,

$$\begin{aligned} v^{cbf} &= \arg \min_{v^{cbf}} \|v^{cbf}\|_2 + K_\epsilon \epsilon \\ \text{s.t. } & A v^{cbf} \leq b, \\ \text{where } & b = [b_1, b_2, \dots, b_k], \\ \text{with } & b_i = p_i^T f(s_t) + p_i^T (\eta - 1) s_t + \\ & \eta q_i + p_i^T g(s_t) v \end{aligned} \quad (15)$$

Note that the above optimisation problem (15) depends on a CBF constraint h of the form (8). A safety constraint for an ego vehicle to maintain a safe time headway from the leading vehicle C_{ol} can be defined as follows

$$h_{ol}(x_t) = \Delta x_{ol} - x^{safe} \quad (16)$$

where Δx_{ol} is the longitudinal distance from the rear end of the preceding vehicle and the front of the ego vehicle as shown in Fig. 2. The x^{safe} is the safe longitudinal distance that must be maintained. For an ego vehicle travelling with a velocity v_e , the safe distance threshold is calculated based on the time headway τ is defined as

$$x^{safe} = \tau * v_e \quad (17)$$

In the above Eq (17), it is important to consider that the safe distance threshold increases as the vehicle speeds up. Consequently, a vehicle may reach an unsafe state even after actuating safe control inputs, as the safe distance threshold at the current step may be higher than in the previous step. To overcome this problem, an additional buffer distance x^{buff} is

added to the safe distance threshold (17) to account for the dynamic change, as shown below

$$x^{\text{safe}} = \tau * v_e + x^{\text{buff}} \quad (18)$$

where, $x^{\text{buff}} = (a^{\text{max}} + 0.1) * \Delta t * \tau$

where Δt is a unit time step and a^{max} is the maximum acceleration of the vehicle. A small constant 0.1 is added to the maximum acceleration a^{max} to restrict a vehicle from moving close to violating the threshold headway distance ($h_{\text{ol}} < 0.001$ m).

2) *Lateral control*: Safe longitudinal distance must be maintained from the adjacent vehicles in the target lane, C_{otl} and C_{otr} , to perform a safe lane change manoeuvre and continue to maintain safety after completing the manoeuvre. Therefore, lane change rules are defined based on the safety invariance condition (7) to define lateral constraints as follows

$$\begin{aligned} h_{\text{otl}}(x_{t+1}) + (\eta - 1)h_{\text{otl}}(x_t) &\geq 0 \\ h_{\text{otr}}(x_{t+1}) + (\eta - 1)h_{\text{otr}}(x_t) &\geq 0 \end{aligned} \quad (19)$$

where the CBF constraints h_* are defined similarly to h_{lon} from Eq. (16) by replacing the longitudinal distance Δx_{ol} with the longitudinal distances, Δx_{otl} and Δx_{otr} , from vehicles C_{otl} and C_{otr} respectively as follows

$$\begin{aligned} h_{\text{otl}} &= \Delta x_{\text{otl}} - x^{\text{safe}} \\ h_{\text{otr}} &= \Delta x_{\text{otr}} - x^{\text{safe}} \end{aligned} \quad (20)$$

The x^{safe} in the h_{otr} is calculated based on the velocity of the rear adjacent vehicle v_{otr} instead of the ego vehicle because the ego vehicle must ensure that a safe headway distance is maintained with respect to a rear adjacent vehicle before initiating a lane change. These constraints are applied to check safe conditions to initiate a lane change, as explained later in Section III-B.

3) *Control limits*: Given that the vehicle control inputs are subject to physical limitations, they must be constrained to avoid generating control inputs that are impossible to actuate. Let's consider the velocity range to be $[v^{\text{min}}, v^{\text{max}}]$. This velocity range is integrated as a constraint in the QP-optimisation problem (15) using the following condition

$$v^{\text{min}} \leq v^{\text{cbf}} + v \leq v^{\text{max}} \quad (21)$$

to ensure that the safe control input is within the physical limits of the vehicle.

B. MARL-HSS

Our MARL lane change controller uses the state of the ego vehicle and the states of observed vehicles as input to evaluate a discrete lane change decision at the behavioural layer. To execute the discrete lane change decision, a feedback controller is used at the motion planning layer to generate low-level control inputs such as acceleration a^{ll} and steering angle ψ^{ll} . The HSS, however, uses velocity v_* as control input in Eq. (13). Therefore, the longitudinal control input for HSS, velocity v_e^{ll} , is derived from the acceleration a^{ll} based on the vehicle model (3) as follows

$$v_e^{\text{ll}} = v_e + a^{\text{ll}} \Delta t \quad (22)$$




Fig. 3. HSS block diagram for longitudinal control

Similarly, the target velocity of the observed vehicle v_o^{ll} is evaluated assuming a worst-case acceleration as a control input. These velocities can be combined to define a low-level control input $v^{\text{ll}} = [v_e^{\text{ll}} \ v_o^{\text{ll}}]$.

The longitudinal safety constraints are integrated into the feedback controller using the v^{ll} in place of v to solve the optimisation problem (15) using QP. As this optimisation problem considers the longitudinal CBF constraint (16), the solution evaluated from QP v^{cbf} is added to v^{ll} , as defined in Eq. (14), to evaluate safe velocity v^{safe} . From a component in v^{safe} corresponding to the ego vehicle v_e^{safe} , the safe control input a^{safe} is evaluated using the vehicle model as follows

$$a^{\text{safe}} = \frac{v_e^{\text{safe}} - v_e}{\Delta t} \quad (23)$$

The block diagram in Fig. 3 summarises the integration of HSS and the motion planning layer to evaluate safe longitudinal control a^{safe} by overriding the possibly unsafe control input a^{ll} from the motion planning layer.

The lateral safety constraints defined in the Eq. (19) use the low-level velocity v^{ll} in place of v to evaluate $h(x_{t+1})$. If the lateral constraints are not satisfied when the behavioural layer decides to change lanes, then a new steering input ψ^{safe} is evaluated using a feedback controller to follow the current lane instead. The lateral safety constraints ensure that the lane change is initiated only when longitudinal safety constraints are maintained with both the leading and following vehicles in the target lane. As these constraints are checked during a lane change, a vehicle moves back towards the centre of the current lane if a lane change is considered unsafe midway during a manoeuvre. Therefore, this check avoids collisions between vehicles moving parallelly the adjacent lanes. Since the lateral safety constraints ensure safety during the lane change and the longitudinal constraints ensure safety before and after the lane change, a vehicle is considered to make safe driving manoeuvres overall.

To integrate conditions to honour physical control limits, the velocity range $[v^{\text{min}}, v^{\text{max}}]$ is evaluated using a vehicle model based on the current velocity v and the acceleration range $[a^{\text{min}}, a^{\text{max}}]$ of the CAVs, as follows

$$\begin{aligned} v^{\text{min}} &= v + a^{\text{min}} \Delta t \\ v^{\text{max}} &= v + a^{\text{max}} \Delta t \end{aligned} \quad (24)$$

This velocity range $[v^{\text{min}}, v^{\text{max}}]$ is used in the Eq. (21) to constrain the control limits of vehicle.

To sum up, an architecture to integrate MARL (Section II-B) and HSS (Section III-A) is presented in Fig. 4. The longitudi-




Fig. 4. MARL-HSS vehicle controller architecture

nal constraints in HSS ensure safe manoeuvres while following a lane before and after the lane change. During a lane change manoeuvre, lateral constraints in HSS ensure safety. Moreover, the control inputs for executing these manoeuvres are constrained to be within the physical limits of CAVs. Since all the CAVs individually evaluate safe control inputs assuming worst-case action from other vehicles, a multi-agent system of CAVs can be considered to be safe [14].

IV. EVALUATIONS AND RESULTS

In this section, the proposed method, MARL-HSS, is compared with unsafe MARL defined in MARL-CAV [11], which is used as the baseline. As mentioned earlier, MARL-CAV is a state-of-the-art lane change controller for making efficient lane decisions to improve traffic efficiency. To evaluate these methods, first a simulation setup is defined for training MARL. Then, the safety guarantee from HSS is demonstrated through minimum time headway observed while training. Next, the learning curves are illustrated to evaluate the training efficiency of MARL-HSS. Finally, MARL policies trained with MARL-HSS and the baseline are compared to analyse the trade-off between safety and efficiency.

A. Evaluation scenario and simulation setup

An on-ramp merging scenario is used to train and evaluate the MARL lane change controllers. This scenario is closely similar to the scenario used in the baseline MARL-CAV [11] with minor changes in the road layout, initial positions, and traffic densities. A road layout of 1420 m is considered in our scenario consisting of 220 m straight road to enter the merging area from the highway and a merging road followed by a 100 m ramp that leads to a merging section of 100 m. After the merging section, the highway road is further extended by 1000 m to allow vehicles to continue moving for 100 steps to complete a training episode. Within the first 320 m, CAVs are spawned at random initial positions separated by 50 m to ensure that CAVs are in a safe state at the start, as




Fig. 5. On-ramp merging scenario

shown in Fig. 5. In this merging scenario, two levels of traffic densities are considered, namely the *Light traffic* consisting of 2-6 CAVs, and the *Moderate traffic* consisting of 4-8 CAVs. The dense traffic with more than 8 CAVs is not considered as it may not create enough opportunities for safe lane changes.

A customised simulator is used to simulate the merging scenario by extending the *highway-env* simulator [23], which is widely used to evaluate AI-based AV controllers. The custom simulator is tailored to integrate the simulation environment with the MARL algorithm (MAPPO) and HSS. The hyperparameters from the baseline [11], such as learning rate, batch size, and reward weights (w_* from (1)), are used to train MARL-HSS for a fair comparison. The policies are trained with 3 random seeds to capture the variability of rewards. During the training for 20,000 episodes, intermediate policies are evaluated using 20 episodes for every 200 training episodes. In the HSS, CBF constraints are configured to maintain a 0.5 s time headway with a leading vehicle, as CAVs are expected to maintain a smaller headway compared to AVs or human-driven vehicles [24]. To define the strictness of CBF constraints, the coefficient of zero order control, used in the Eq. (9), is configured as $\eta = 0.0325$. This value is empirically chosen using binary search to ensure safety compliance and smooth transitions in safe control inputs evaluated from HSS. In this simulation, the behavioural layers make control decisions at a 5 Hz frequency, and the motion planning layers operate at a 15 Hz frequency. Therefore, MARL must evaluate the lane change decision within 0.2 s. Moreover, the motion planning layer consisting of the feedback controller and HSS must evaluate the safe control action within 0.0667 s. Both the baseline and MARL-HSS lane change controllers were trained in a shared Ubuntu 24.04 server with AMD EPYC 9654 processor, 4×40 GB GPUs (Nvidia-L40), and 1.5 TB memory. They only used, however, approximately 1.5 GB GPU memory and 3 GB RAM memory.

B. Safety evaluation

In this subsection, the safety guarantee from HSS is analysed in moderately dense traffic that offers challenging lane change scenarios with limited opportunities for safe lane changes. In such scenarios, MARL policy will push the CAVs towards violating the time headway safety constraint (0.5 s) to achieve higher rewards. Typically, time headway is a positive number indicating a time gap between two vehicles. However,




Fig. 6. Minimum time headway observed during the MARL training in moderate traffic. The shaded region illustrates the range of minimum time headway observed over 3 random seeds.

when a collision occurs, negative values are observed in the simulation. The minimum time headway is captured by checking the time headway of each CAV at each step of the episodes used for the intermediate evaluations while training. The minimum time headways observed with unsafe MARL and MARL-HSS are compared in Fig. 6. With the baseline controllers, CAVs frequently violate the safety constraint, leading to crashes. These violations can be avoided with HSS integration as MARL-HSS does not cross the safety constraint of 0.5 s time headway. The proposed safety shield, therefore, ensures that the dynamic safety constraint defined using time headway is enforced strictly.

C. MARL training

The reward curves from MARL-HSS are compared with the baseline in Fig. 7 to evaluate the training performance. The reward curves illustrate a higher sample efficiency of the proposed method compared to the baseline. As light traffic provides sufficient opportunities for safe lane changes, our approach comfortably achieves higher rewards after 13,200 episodes of training, as shown in Fig. 7a. Without a safety shield, the unsafe MARL needs to explore more episodes compared to MARL-HSS to achieve higher rewards. Therefore, it achieves higher rewards towards the later stages of training after 19,800 episodes. Similarly, while the proposed method learns the best policy after 16,200 episodes, the baseline requires 18,400 episodes of training, as illustrated in Fig. 7b.

Importantly, the proposed method learns a policy that is more stable than the baseline, as shown by unsafe MARL policies having higher standard errors than MARL-HSS for both traffic levels. Moreover, our method achieves more consistent rewards throughout the training process in both traffic levels compared to the unsafe MARL's rewards that drop noticeably.

D. MARL policy evaluation

In this subsection, the best-performing policies are evaluated using 100 *policy evaluation episodes*, which are different from




Fig. 7. Comparison of the reward curves during trainings in *Light* and *Moderate* traffic levels. The spread indicates the standard error observed over 3 random seeds. The curves are smoothed using an average over a window of 5 evaluation epochs

TABLE I
POLICY EVALUATION OF MARL-HSS AND UNSAFE MARL IN DIFFERENT TRAFFIC LEVELS USING 100 EPISODES

	Traffic level	Training episodes	Speed (m/s)	Crash count
MARL-HSS	Light	13,200	28.36	0
	Moderate	16,200	26.54	0
Unsafe MARL	Light	19,800	27.70	6.3
	Moderate	18,400	27.32	5

the episodes used for intermediate evaluation while training. The policies are compared on the basis of speed (m/s) to evaluate traffic efficiency, and crash counts are considered for safety evaluation. The average values of these parameters over the policy evaluation episodes are presented in Table I.

The results in Table I show that MARL-HSS balances a trade-off between safety and efficiency. With the MARL-HSS policy, zero crashes are observed in both traffic levels, which reaffirms safety guarantees. However, mean crash counts of 6.3 and 5 are observed with unsafe MARL in light and moderate traffic, respectively. In a light traffic density, the policy with the proposed safety shield marginally improves the traffic efficiency with a 2.38% increase in the average speed. A comparable traffic efficiency is observed between the proposed method and the baseline in moderate traffic with a marginal difference of 0.78 m/s in the average speeds. With closely similar average speeds, MARL-HSS is considered to achieve traffic efficiency of the baseline while respecting safety constraints.

V. CONCLUSION

In this article, we propose a decentralised safety shield, HSS, to ensure safety based on dynamic constraints. This safety shield formulates CBF constraints for longitudinal and

lateral control to override possibly unsafe control inputs with a minimum correction to ensure safe manoeuvres. To integrate HSS into the MARL lane change controller, we define the MARL-HSS architecture. In this architecture, the behavioural control decisions from the lane change controller are translated into safe control inputs at the motion planning layer. Experimental evaluation shows that the dynamic safety constraints are strictly enforced, preventing unsafe actions even when MARL prioritises higher rewards. Moreover, the proposed lane change controller learns a stable policy with less variability in the training. While the policy trained with MARL-HSS outperforms the state-of-the-art baseline in light traffic density, it achieves comparable traffic efficiency in moderate traffic. Notably, MARL-HSS balances a trade-off between safety and traffic efficiency by ensuring safety across different traffic densities without significantly compromising on the average speed.

MARL-HSS could be extended in the future to support mixed traffic, by considering uncertainties associated with vehicles of variable levels of autonomy and connectivity. Moreover, exploiting CAV's ability to collaboratively decide control actions in the safety shield could potentially further improve traffic efficiency while ensuring safety.

ACKNOWLEDGMENT

The authors wish to thank the editors and anonymous reviewers for their helpful suggestions. This publication has emanated from research conducted with the financial support of Taighde Éireann – Research Ireland under Grant numbers 18/CRT/6222 at Centre for Research Training in Advanced Networks for Sustainable Societies (ADVANCE CRT) and 13/RC/2077_P2 at CONNECT: the Research Ireland Centre for Future Networks.

REFERENCES

- [1] S. Agrawal, A. M. Schuster, N. Britt, E. A. Mack, M. L. Tidwell, and S. R. Cotten, "Building on the past to help prepare the workforce for the future with automated vehicles: A systematic review of automated passenger vehicle deployment timelines," *Technology in Society*, vol. 72, p. 102186, Feb. 2023.
- [2] Y. Ma, Z. Wang, H. Yang, and L. Yang, "Artificial intelligence applications in the development of autonomous vehicles: a survey," *IEEE/CAA Journal of Automatica Sinica*, vol. 7, no. 2, pp. 315–329, Mar. 2020.
- [3] D. Marinescu, J. Čurn, M. Bourroche, and V. Cahill, "On-ramp traffic merging using cooperative intelligent vehicles: A slot-based approach," in *2012 15th International IEEE Conference on Intelligent Transportation Systems*, Sep. 2012, pp. 900–906, ISSN: 2153-0017.
- [4] B. Hegde and M. Bourroche, "Design of AI-based lane changing modules in connected and autonomous vehicles: a survey," in *Twelfth International Workshop on Agents in Traffic and Transportation*, Vienna, 2022.
- [5] W. Liu, M. Hua, Z. Deng, Z. Meng, Y. Huang, C. Hu, S. Song, L. Gao, C. Liu, B. Shuai, A. Khajepour, L. Xiong, and X. Xia, "A Systematic Survey of Control Techniques and Applications in Connected and Automated Vehicles," *IEEE Internet of Things Journal*, vol. 10, no. 24, pp. 21 892–21 916, Dec. 2023, conference Name: IEEE Internet of Things Journal.
- [6] C. Yu, X. Wang, X. Xu, M. Zhang, H. Ge, J. Ren, L. Sun, B. Chen, and G. Tan, "Distributed Multiagent Coordinated Learning for Autonomous Driving in Highways Based on Dynamic Coordination Graphs," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 2, pp. 735–748, Feb. 2020.
- [7] J. Dong, S. Chen, Y. Li, R. Du, A. Steinfeld, and S. Labi, "Space-weighted information fusion using deep reinforcement learning: The context of tactical control of lane-changing autonomous vehicles and connectivity range assessment," *Transportation Research Part C: Emerging Technologies*, vol. 128, p. 103192, Jul. 2021.
- [8] S. Chen, J. Dong, P. Y. J. Ha, Y. Li, and S. Labi, "Graph neural network and reinforcement learning for multi-agent cooperative control of connected autonomous vehicles," *Computer-Aided Civil and Infrastructure Engineering*, vol. 36, no. 7, pp. 838–857, 2021.
- [9] W. Zhou, D. Chen, J. Yan, Z. Li, H. Yin, and W. Ge, "Multi-agent reinforcement learning for cooperative lane changing of connected and autonomous vehicles in mixed traffic," *Autonomous Intelligent Systems*, vol. 2, no. 1, p. 5, Mar. 2022.
- [10] Y. Hou and P. Graf, "Decentralized Cooperative Lane Changing at Free-way Weaving Areas Using Multi-Agent Deep Reinforcement Learning," *arXiv:2110.08124 [cs]*, Oct. 2021, arXiv: 2110.08124.
- [11] D. Chen, M. R. Hajidavalloo, Z. Li, K. Chen, Y. Wang, L. Jiang, and Y. Wang, "Deep Multi-Agent Reinforcement Learning for Highway On-Ramp Merging in Mixed Traffic," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–16, 2023.
- [12] B. Hegde and M. Bourroche, "Multi-agent reinforcement learning for safe lane changes by connected and autonomous vehicles: A survey," *AI Communications*, vol. 37, no. 2, pp. 203–222, Jan. 2024, publisher: IOS Press.
- [13] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, "End-to-End Safe Reinforcement Learning through Barrier Functions for Safety-Critical Continuous Control Tasks," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, pp. 3387–3395, Jul. 2019, number: 01.
- [14] S. Han, S. Zhou, J. Wang, L. Pepin, C. Ding, J. Fu, and F. Miao, "A Multi-Agent Reinforcement Learning Approach for Safe and Efficient Behavior Planning of Connected Autonomous Vehicles," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–17, 2023, conference Name: IEEE Transactions on Intelligent Transportation Systems.
- [15] X. Wang, "Ensuring Safety of Learning-Based Motion Planners Using Control Barrier Functions," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4773–4780, Apr. 2022.
- [16] M. Garg and M. Bourroche, "Can Connected Autonomous Vehicles Improve Mixed Traffic Safety Without Compromising Efficiency in Realistic Scenarios?" *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 6, pp. 6674–6689, Jun. 2023.
- [17] B. Hegde and M. Bourroche, "Safe CAV lane changes using MARL and control barrier functions," in *Thirteenth International Workshop on Agents in Traffic and Transportation*, vol. 3813, Santiago de Compostela, Spain, Oct. 2024, pp. 83–95.
- [18] B. Paden, M. Cáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A Survey of Motion Planning and Control Techniques for Self-Driving Urban Vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, Mar. 2016.
- [19] A. Matin and H. Dia, "Impacts of Connected and Automated Vehicles on Road Safety and Efficiency: A Systematic Literature Review," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–0, 2022.
- [20] P. Polack, F. Althché, B. d'Andréa Novel, and A. de La Fortelle, "The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles?" in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 812–818.
- [21] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control Barrier Functions: Theory and Applications," in *2019 18th European Control Conference (ECC)*, Jun. 2019, pp. 3420–3431.
- [22] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge New York Melbourne New Delhi Singapore: Cambridge University Press, 2004.
- [23] E. Leurent, "An Environment for Autonomous Driving Decision-Making," 2018, publication Title: GitHub repository. [Online]. Available: <https://github.com/eleurent/highway-env>
- [24] M. Garg, C. Johnston, and M. Bourroche, "Can Connected Autonomous Vehicles really improve mixed traffic efficiency in realistic scenarios?" in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, Sep. 2021, pp. 2011–2018.