Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnline CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 55-14 (2022) 1–6

Safe eco-cruise control for connected vehicles against trained cyber attacks ⋆

András Mihály * **Balázs Németh** * **Péter Gáspár** *

* *Systems and Control Laboratory, Institute for Computer Science and Control (SZTAKI), Eötvös Loránd Research Network (ELKH), Kende u. 13-17, H-1111 Budapest, Hungary. E-mail: [andras.mihaly;balazs.nemeth;peter.gaspar]@sztaki.hu*

Abstract:
The paper proposes an adaptive cruise control method for connected and automated vehicles (CAVs) with safety considerations against cyber attacks. A high-level layer is responsible for the computation of energy optimal speed profiles for the CAVs, considering oncoming road information such as terrain characteristics and speed limits. Due to the computationally cumbersome optimization method of the speed profile design, this step is performed in a cloud. Next, a feasibility analysis is carried out on the vehicle layer regarding safety of the CAVs, overwriting high-level speed references in case of a collision risk is detected. The aim of the present paper is to validate the above multi-layer control method with the design of an intelligent cyber attack using reinforcement learning techniques. Evaluating a multi-agent training with real data velocity profiles, each automated vehicle has been simulated to be attacked by an agent aiming to generate collisions in the vehicle string.

Keywords: cyber attack, reinforcement learning, cruise control, connected vehicles.

1. INTRODUCTION AND MOTIVATION

One of the most important aspects in automated driving scenarios is the reduction of energy consumption, which require the use of vehicle navigation systems, perception sensors and intelligent design methods, see Sciarretta and Vahidi (2019). The aim of eco-driving is to enhance energy efficiency by adapting the speed of the automated vehicle to the environment, taking into consideration terrain characteristics, speed limits and traffic flow as well as the behavior of the surrounding vehicles. A predictive robust and optimal control method has been introduced in Gáspár and Németh (2019) which is used as the baseline controller in present paper.

Note, that in the literature several other eco-cruise control methods have been introduced, with the majority of them using classical optimization-based solutions, see Padilla et al. (2018); Passenberg et al. (2009); Hellström et al. (2009); Saerens et al. (2013). Moreover, recently learning-based eco-cruise control methods have been presented using Q-learning algorithms, artificial neural networks and deep learning-based solutions, see Bougiouklis et al. (2018); Zhu et al. (2019); Liu et al. (2017); Wu et al. (2019). Although all of the above listed methods may give an optimal solution, the big computational effort required for solving the multi-criteria optimization task can lead to difficulties for on-board vehicle applications. A solution for this problem can be to separate the control design into layers based on functionality. The cumbersome optimization task related to eco-driving can be performed in a cloud, sending the optimal speed reference for the CAVs via internet communication. Hence, information from the cloud must be analyzed before application in the automated vehicle since in case of a cyber attack these signals might be corrupted, see Guo et al. (2019).

A comprehensive survey on the latest results for cyber attacks and defenses for autonomous vehicles has been presented in Kim et al. (2021). Here, attacks are classified into the categories of autonomous control and driving systems, and vehicle-to-everything (V2X) communications. It has been highlighted, that while earlier attack and defense strategies have been conducted on the vehicle CAN and ECU, in recent years attacks on external communications have been in the focus. A physics-guided machine learning method for the detection of cyber attacks on electric vehicle driveline has been introduced in Guo et al. (2021), which demonstrated high accuracy in a hardware-in-the-loop (HIL) simulation testbed. The influence of cyber attacks for CAVs with longitudinal controllers has been analyzed by Li et al. (2018), while a radar sensor health monitoring method using observer for CAVs has been proposed by Jeon et al. (2021). In Dong et al. (2020), the impact of cyber attacks on CAVs has been analyzed related to the traffic flow, including the risk of rear-end collision.

To avoid risks of cyber attacks on CAVs, a method of analyzing vehicle speed has already been introduced earlier in Németh et al. (2021). Here, a cyber attack has been simulated with randomly selected corrupted reference velocities. The aim of the present paper is to vali-

⋆ The research was supported by the Ministry of Innovation and Technology NRDI Office within the framework of the Autonomous Systems National Laboratory Program. The research was partially supported by the TKP2021-NKTA-01 NRDIO grant on "Research on cooperative production and logistics systems to support a competitive and sustainable economy".

10.1016/j.ifacol.2022.07.574

date the multi-layer eco-cruise control method for CAVs by designing a hostile cyber attack using reinforcement learning techniques. Hence, a more expert cyber attack is designed with a systematic attack approach. The training process of the cyber attack agents have been performed on several different velocity profiles based on real-world highway measurement datasets of Next Generation Simulation (NGSIM) with the aim to cause accidents among the CAVs. The primary safety performance guaranteed by the vehicle level speed analysis is to avoid rear-end collision among CAVs, while maintaining a safe following distance. The performed simulations demonstrate both the effectiveness of the designed cyber attack without the guaranteed primary safety performances, as well as the desired operation of the latter vehicle level speed analysis layer.

The paper is organized as follows. The hierarchical multi-layer structure is proposed in Section 2, describing the velocity optimization method and design of the safety layer on the vehicle level. Section 3 introduces the reinforcement learning method of the cyber attack. In Section 4 real data simulations are performed for CAVs under cyber attack both with and without applying the safety layer on the vehicle level. Finally, contributions of the paper are summarized in Section 5.

2. MULTI-LAYER DESIGN FOR THE ECO-CRUISE CONTROL

The eco-cruise control for the CAVs contains three layers, as depicted in Figure 1. The high-level is responsible for the optimization of the vehicle velocity v_{high}, by which the secondary performances as energy efficiency and traveling time can be met. This cumbersome calculation is practically computed in a cloud, and the speed reference signal is transmitted to the vehicle through internet communication. In the vehicle level, a speed analysis layer is responsible to guarantee the primary safety related performances based on the measurements of the onboard sensors. In case the primary performances are not violated by applying the high level velocity for the vehicle, it is forwarded to the vehicle control layer as $v_{veh} = v_{high}$. On the other hand, in case the safety performances can not be guaranteed with v_{high}, it is necessary to compute a modified v_{veh} reference signal for the vehicle control. Finally, the third layer contains the local speed controller of the vehicle, responsible for tracking v_{veh} by applying adequate driveline actuation.

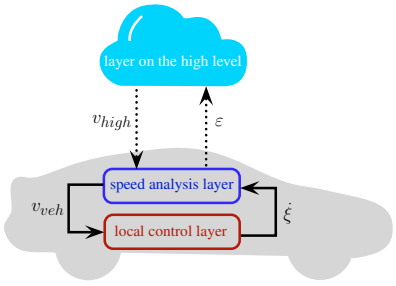

Fig. 1. Architecture of the eco-cruise control

2.1 Velocity design considering secondary performances

As the look-ahead control method for velocity design has already been detailed in Gáspár and Németh (2019); Németh and Gáspár (2017), here only a brief summary is given. It is assumed, that the road ahead of the vehicle is divided into n number of segments with corresponding reference velocities $v_{ref,i}$, $i \in \{1,..,n\}$, which are the speed limits. Next, prediction weights Q, γ_i, $i \in \{1,..,n\}$ are given for the road segments defining their importance in the speed design. Note, that while Q defines the tracking of the actual reference speed $v_{ref,0}$, γ_i weights stands for the consideration of further road slopes and speed limits.

The speed given by the look-ahead optimization algorithm guarantees an optimal balance between energy consumption and traveling time, as the oncoming road slopes and speed limits are considered in the design. Road slopes given by α_i are incorporated in the longitudinal force resistances $F_{di,r}$, while speed limits v_i are also given for the look-ahead road sections. With the former considerations, the following formula is derived for the optimal velocity: $v_{high} = \sqrt{\vartheta - 2s_1(1-Q)(\ddot{\xi}_0 + gsin\alpha)}$, where $\ddot{\xi}_0$ is the vehicle acceleration, s_1 is the distance of the actual road section, α is the actual road slope and ϑ contains the look-ahead information:

$$\vartheta = Qv_{ref,0}^2 + \sum_{i=1}^{n} \gamma_i v_{ref,i}^2 + \frac{2}{m} \sum_{i=1}^{n} s_i F_{di,r} \sum_{j=i}^{n} \gamma_j. \quad (1)$$

To ensure an optimal velocity for the vehicle, the longitudinal control force and the traveling time must be minimized at the same time. For the former criterion $F_{l1}^2 \rightarrow \min$, a quadratic optimization is solved by selecting weights:

$$\bar{F}_{l1}^2 = (\beta_0(\bar{Q}) + \beta_1(\bar{Q})\bar{\gamma}_1 + \ldots + \beta_n(\bar{Q})\bar{\gamma}_n)^2 \rightarrow \min \quad (2)$$

with constraints $0 \leq \bar{Q}, \bar{\gamma}_i \leq 1$ and $\bar{Q} + \sum \bar{\gamma}_i = 1$. The traveling time criteria requires to minimize the difference between the actual velocity and the speed limit, which leads to the optimization problem $|v_{ref,0} - \dot{\xi}_0| \rightarrow \min$, whose solution is achieved by choosing the weights $\breve{Q} = 1$ and $\breve{\gamma}_i = 0, i \in [1, n]$.

The balance between the minimization of longitudinal force and traveling time is given by prediction weights as follows:

$$Q = R_1\bar{Q} + R_2\breve{Q} = 1 - R_1(1-\bar{Q}) \quad (3a)$$
$$\gamma_i = R_1\bar{\gamma}_i + R_2\breve{\gamma}_i = R_1\bar{\gamma}_i, \quad i \in \{1,..,n\} \quad (3b)$$

where $\bar{Q}, \bar{\gamma}_i$ are the energy-optimal, while $\breve{Q}, \breve{\gamma}_i$ are the time-optimal weights.

Note, that weight $R_1 \in [0;1]$ in (3) is responsible for creating the balance between the defined criteria, with $R_1 = 1$ leading to low energy consumption and more varying speed profile, whereas with $R_1 = 0$ the values of the speed limit are selected for the vehicle to follow.

The R_1 selection method is founded on the performances of the eco-cruise control problem, such as the performance of speed variation between limitations and the performance of energy consumption minimization. It is determined by the bounds of the reference speed variation as

$$v_{ref,0} + \Delta_{l,min} \leq v_{high} \leq v_{ref,0} + \Delta_{l,max}, \quad (4)$$

where $\Delta_{l,max}, \Delta_{l,min}$ are predefined scalars defining the bounds of the reference speed variation. Hence, to ensure constraint (4), the following optimization problem is formulated:

$$\max R_1 \tag{5a}$$

subject to

$$v_{ref,0} + \Delta_{l,min} \leq v_{high} \leq v_{ref,0} + \Delta_{l,max}. \tag{5b}$$

Finally, the computation of v_{high} is given as follows: with $v_{ref,i}, \alpha_i$ the weights $\bar{Q}, \bar{\gamma}_i$ are calculated applying the quadratic optimization (2). Next, R_1 value must be found, which ensures the predefined constraint given in (5).

2.2 Safety performances in the cruise control design

The speed analysis guaranteeing the primary performances has already been detailed in Németh et al. (2021), thus here a brief summary of the method is given. The goal of the speed analysis is to calculate a reference velocity v_{veh} for the local controller, by which the safety performances can be ensured. The calculation of v_{veh} is founded on the analysis of v_{high}, which is sent by the high level layer via internet communication.

The analysis is designed to check, whether the following conditions hold:

(1) The CAVs must maintain a safe distance d_{safe} from the preceding vehicle. This condition results in a maximum reference speed.
(2) The CAVs must maintain a safe distance d_{safe} from the follower vehicle. This condition results in a minimum reference speed.
(3) In case their is no other vehicle present in the environment, the velocity must be set between $v_{ref,0}+\Delta_{l,min}$ and $v_{ref,0} + \Delta_{l,max}$.

The analysis of these conditions requires the prediction of the forthcoming distances from the preceding and follower vehicles. Guaranteeing safety performance requirement $d_{prec}(T) \leq d_{safe}$ yields the following relation:

$$d_{safe} \leq \frac{\ddot{\eta}^p(0)T^2}{2} - \frac{\ddot{\xi}(0)T^2}{2} + \dot{\eta}^p(0)T - \dot{\xi}(0)T + d_{prec}(0), \tag{6}$$

where 0 represents the actual time, T is the prediction time horizon, d_{prec} is the forthcoming distance between the preceding vehicle and the ego vehicle, $d^p_{prec}(0)$ is the actual distance, while $\ddot{\eta}^p(0)$ is the acceleration, $\dot{\eta}^p(0)$ is the speed of the preceding vehicle.

While $\dot{\eta}^p(0), d_{prec}(0)$ can be measured easily using camera or radar, the measurement of $\ddot{\eta}^p(0)$ may contain significant noise, thus it is over-approximated by a constant value $a_{min} \leq \ddot{\eta}^p(0)$ representing maximum braking. Thus, (6) can be reformulated as follows:

$$d_{safe} \leq \frac{a_{min}T^2}{2} - \frac{\ddot{\xi}(0)T^2}{2} + \dot{\eta}^p(0)T - \dot{\xi}(0)T + d_{prec}(0) \tag{7}$$

and thus, the acceleration of the vehicle must be limited as

$$\ddot{\xi}(0) \leq a_{min} + \frac{2\dot{\eta}^p(0)}{T} - \frac{2\dot{\xi}(0)}{T} + 2\frac{d_{prec}(0) - d_{safe}}{T^2}. \tag{8}$$

Assuming constant acceleration $\ddot{\xi}(0) = \frac{\dot{\xi}(T)-\dot{\xi}(0)}{T}$ on the prediction horizon, the required speed of the vehicle in T is computed as $\dot{\xi}(T) = \dot{\xi}(0) + \ddot{\xi}(0)T$. Hence, the maximum of the reference speed is given as follows:

$$v_{veh} \leq a_{min}T + 2\dot{\eta}^p(0) - \dot{\xi}(0) \quad +2\frac{d_{prec}(0) - d_{safe}}{T}. \tag{9}$$

Similarly, the predicted distance among the ego vehicle and the follower vehicle d_{foll} can be calculated from the second derivative of the distance between them as $\ddot{d}_{foll} = \ddot{\xi} - \eta^f$. Applying a maximum acceleration a_{max} for the follower vehicle with expression $a_{max} \geq \ddot{\eta}^f(0)$, the minimum reference speed is given as follows:

$$v_{veh} \geq \dot{\xi}(0) - a_{max}T + 2\dot{\eta}^f(0) \quad +2\frac{d_{safe} - d_{foll}(0)}{T}. \tag{10}$$

The design of the reference speed value v_{veh} for the vehicle is thus given as follows. In case $v_{veh} = v_{high}$ satisfies the inequalities (9) and (10), v_{high} can be transmitted directly to the local speed controller.

The third primary performance requirement is to keep vehicle speed in a limited range of $v_{ref,0}$. This information can be extracted from camera-based sign recognition system Bangquan and Xiao Xiong (2019), a digital map Liu et al. (2020), or the fusion of the two systems. As the high level layer ensures that v_{high} is inside of the requested velocity range, in case of $v_{veh} = v_{high}$ the third primary performance requirement is guaranteed.

While the value of v_{high} can be given for v_{veh}, in some cases one of the inequalities (9)-(10) is violated. In these instants an appropriate value for v_{veh} must be selected close as possible to v_{high}, such that the primary performances are ensured at the same time. Hence, if v_{high} fails to meet the inequalities (9)-(10) the following optimization algorithm is performed:

$$\min_{v_{veh}} \ |v_{veh} - v_{high}| \tag{11a}$$

subject to

$$v_{veh} \leq a_{min}T + 2\dot{\eta}^p(0) - \dot{\xi}(0) + 2\frac{d_{prec}(0) - d_{safe}}{T}, \tag{11b}$$

$$v_{veh} \geq \dot{\xi}(0) - a_{max}T + 2\dot{\eta}^f(0) + 2\frac{d_{safe} - d_{prec}(0)}{T}, \tag{11c}$$

$$v_{ref,0} + \Delta_{l,min} \leq v_{veh} \leq v_{ref,0} + \Delta_{l,max}. \tag{11d}$$

3. DESIGN OF CYBER ATTACK WITH REINFORCEMENT LEARNING

As detailed earlier, the main goal of the proposed hierarchical structure for CAVs is to handle a hostile cyber attack by preserving safe motion of vehicles, moreover guaranteeing collision avoidance. In order to validate the effectiveness of the proposed method detailed in Section 2.2, a cyber attack has been designed and applied on the CAVs as shown in Figure 2.

In this scheme, it is assumed that each vehicle obtains an optimal velocity v_{high} calculated in the cloud based on available road data, which they have to follow to meet the

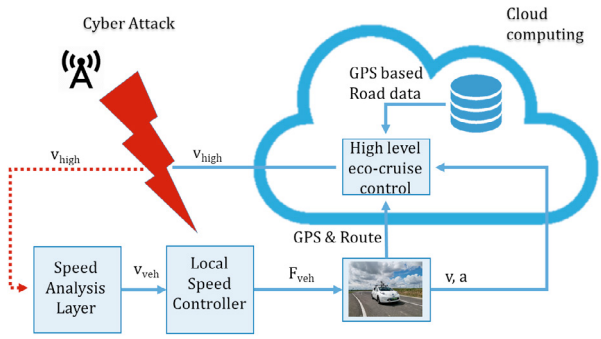

Fig. 2. Scheme of the proposed multi-layer controller considering cyber attacks

requirements of the eco-driving features. However, since the prescribed target velocity for the autonomous vehicles are sent through communication channels which can be vulnerable to cyber attacks, the values of v_{high} can be corrupted. The aim of the cyber attack design is to select a velocity v_{high} for each CAV in a manner, by which the possibility of collisions can be maximized. Hence, for this purpose reinforcement learning procedure has been developed based on the dynamic model of a vehicle platoon, where the vehicle agents have been trained to select a high-level velocity reference for each CAV to follow in order to maximize the probability of a collision. Multiple deep deterministic policy gradient (DDPG) agents have been trained during simulation for three ego vehicles following the lead vehicle, as depicted in Figure 3. Agents have been trained to achieve collision with the speed analysis layer built in the simulation for a more effective protection for the vehicle platoon.

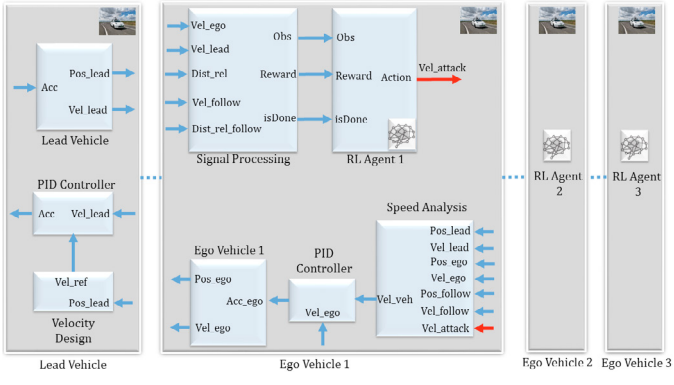

Fig. 3. Cyber attack reinforcement learning procedure with speed analysis

The lead vehicle follows a predefined velocity profile given by the eco-cruise control system or based on real measurement data of the Federal Highway Administration Research and Technology, U.S. Department of Transportation. For the latter the NGSIM computer program collected detailed, high-quality traffic datasets on the road southbound US 101, also known as the Hollywood Freeway, in Los Angeles. The three agents have been trained to control the velocity of the vehicles in a manner to generate catch-up collision in the vehicle string, which the speed analysis layer is designed to prevent. The training models for the ego vehicles has been set as follows:

- The velocity action signal from the agent to the environment (vehicle) is from 0 to 50 m/s.
- The observations from the environment are the following:
  - Ego vehicle actual velocity: $\dot{\xi}(t)$
  - Velocity error from preceding vehicle: $\dot{\eta^p}(t) - \dot{\xi}(t)$
  - Integral of the velocity error from preceding vehicle: $\int(\dot{\eta^p}(t) - \dot{\xi}(t))dt$
  - Velocity error from following vehicle: $\dot{\xi}(t) - \dot{\eta^f}(t)$
  - Integral of the velocity error from following vehicle: $\int(\dot{\xi}(t) - \dot{\eta^f}(t))dt$
  - Following vehicle actual velocity: $\dot{\eta^f}(t)$

The reward r_t provided at every time step t, is $r_t = \min(d^p_{rel}, d^f_{rel})$, where d^p_{rel} is the relative distance from the preceding vehicle, while d^f_{rel} is the relative distance from the following vehicle. Note, that the episode reward considered during training is the cumulative value of r_t. An example of the training process shown in Figure 3 regarding episode rewards for the three agents are depicted in Figure 4.

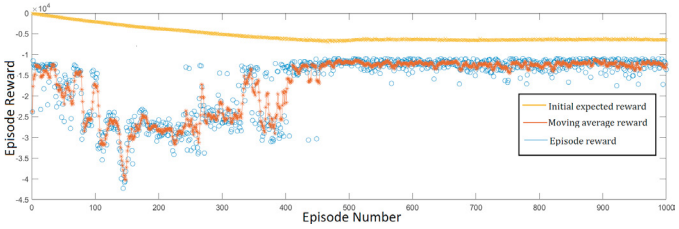

(a) RL Agent 1

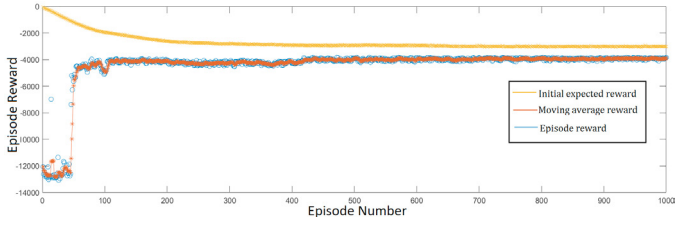

(b) RL Agent 2

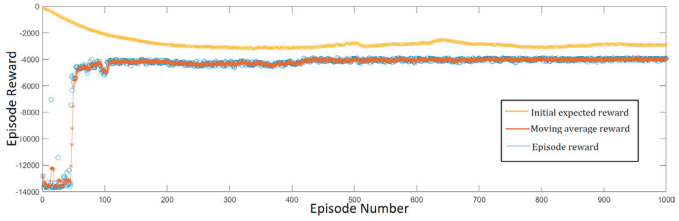

(c) RL Agent 3

Fig. 4. Episode rewards of reinforcement learning agents

Initial positions and velocities have been defined for the vehicles along with physical limitations of the vehicle dynamics, while the sample time of the simulation has been set to $T_s = 0.1\ s$. DDPG agents have been created separately for the three follower vehicle using specified deterministic policy actor and Q-value function critic representation and agent options. The training for the three agents has been set by the maximum episode number and a stopping condition regarding the episode reward value. Note, by separating each vehicle agents, the designed attack can be more specific for the location of the vehicle in the

string, i.e. an agent is trained for the follower of the leader, for the vehicle in between CAVs, and for the last vehicle in the string without follower. During training, a DDPG agent updates the actor and critic properties at each time step and stores past experiences using a circular experience buffer, while action is chosen by the policy using a stochastic noise model at each training step. Note, that the trainings are terminated when the relative distance between CAVs become less than 0.

4. SIMULATION RESULTS

In the first case the simulation contains four vehicles which travel on a highway section. The leader vehicle follows a velocity profile given by the eco-cruise control system, while the following vehicles high-level velocity profiles are under a cyber attack from the beginning of the journey. Note, that the three agents have been previously trained as depicted in Figure 3 for more than 1600 episodes. Next, simulations have been evaluated without and with the low-level speed analysis detailed in Section 2.2 using $a_{min} = -3\ m/s^2$, $a_{max} = 2\ m/s^2$ and $T = 0.5\ s$ as parameters of the algorithm. Note, that in the presented method safe spacing among vehicles are calculated as a function of the actual vehicle velocities as $d_{safe} = \max(d_{min}, \dot{\xi} \cdot t_{safe})$ with using $t_{safe} = 2\ s$ and $d_{min} = 3\ m$ in the simulation. It is well demonstrated in Figure 5 (a), that under the cyber attack without guaranteed performances, the velocities of the autonomous vehicles increase significantly due to the hostile reference high-level speed signals depicted in Figure 5 (b). Hence, the second autonomous vehicle collide with the leader vehicle under less than 10 seconds by catching up, as demonstrated in Figure 5 (c) and Figure 5 (d).

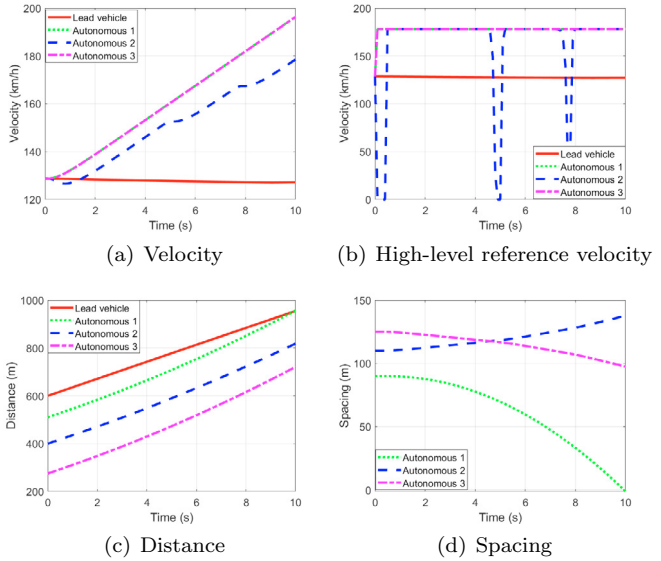

(a) Velocity (b) High-level reference velocity

(c) Distance (d) Spacing

Fig. 5. Cyber attack results without guaranteed performances

Next, simulation was performed with the same initial conditions and trained cyber attack agents applying the low-level speed analysis layer for the autonomous vehicles. It is well demonstrated in Figure 6 (a), that velocity of the autonomous vehicle string adapts to the leader vehicle speed, despite the corrupted high-level velocity signals given by the agents depicted in Figure 6 (b). Hence, collision of vehicles are successfully avoided, as shown in Figure 6 (c). Note, that with lower velocities spacing among autonomous vehicles becomes smaller, as shown in Figure 6 (d).

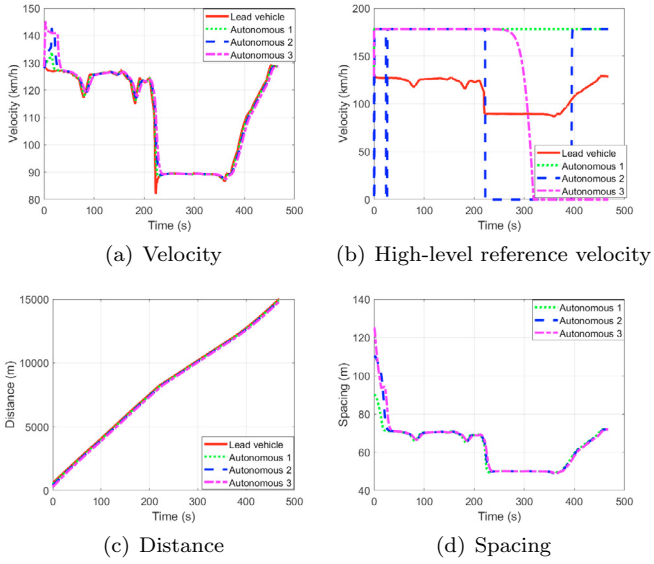

(a) Velocity (b) High-level reference velocity

(c) Distance (d) Spacing

Fig. 6. Cyber attack results with guaranteed performances

In the second case, several simulations have been performed on the real datasets provided by the U.S. Department of Transportation, based on the collected data of a 1 km section of the Hollywood Freeway in Los Angeles. Thus, the similar four vehicle simulation has been performed with the leader vehicle following velocity profiles collected by the NGSIM computer. Note, that for the more uneven velocity profiles the low-level speed analysis has been re-calibrated using $a_{min} = -10\ m/s^2$, $a_{max} = 10\ m/s^2$, while the discrete time PID cruise controller has also been tuned to deal with the bigger accelerations and lower speeds of the leader vehicle. The three agents for the cyber attack have been trained as depicted in Figure 3 for 1000 episodes.

In Figure 7 the results of the cyber attack are depicted without the low-level speed analysis for one of the low speed leader velocity profile. It can be seen in Figure 7 (a) and Figure 7 (b), that without the proposed control layer the first and third autonomous vehicle abruptly increase their velocity, while the second autonomous vehicle decreases its speed. Thus, in less than five seconds one collision occurs between the leader vehicle and the first autonomous vehicle as shown in Figure 7 (c) and Figure 7 (d), while the third and fourth autonomous vehicle would also collide if the simulation had not stopped.

By applying the proposed low-level controller, the reference velocity given by the cyber attack in Figure 8 (b) is overridden as shown in Figure 8 (a). Hence, the second and the third autonomous vehicles comes to a stop around 10 seconds with the 3 meters safety distance preserved, as demonstrated in Figure 8 (c) and Figure 8 (d). Note, that the first autonomous vehicle continues to follow the leader vehicle also preserving the calculated safety distance.

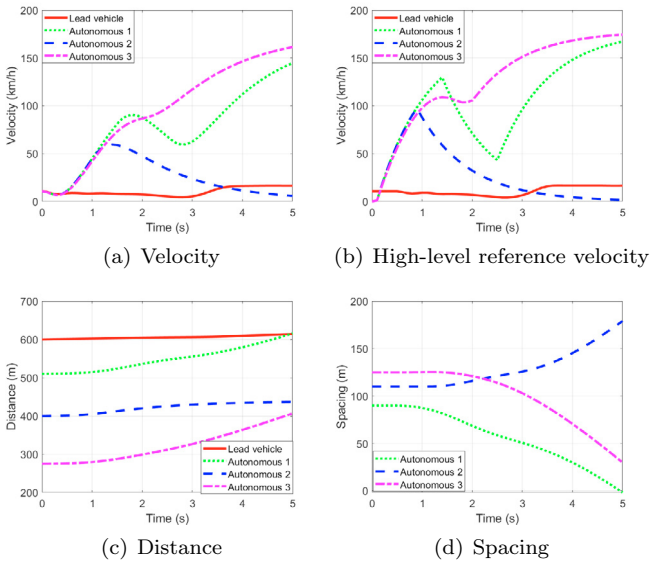

(a) Velocity (b) High-level reference velocity

(c) Distance (d) Spacing

Fig. 7. Cyber attack results without guaranteed performances

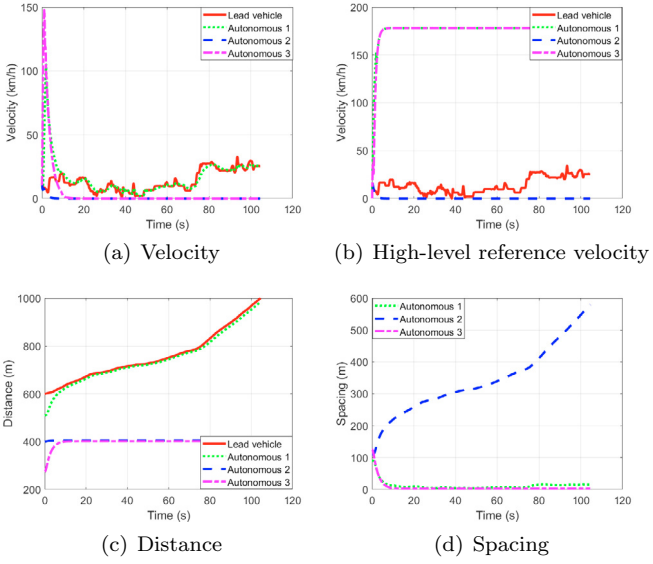

(a) Velocity (b) High-level reference velocity

(c) Distance (d) Spacing

Fig. 8. Cyber attack results with guaranteed performances

5. CONCLUSIONS

The paper showed a strategy for the design of a fault-tolerant cruise control system for CAVs, which is able to handle both faults in the high-level controller or a cyber attack against the communication system of the vehicles. The design has been evaluated in a hierarchical framework with layers related to the cloud computation and the vehicle level. Next, an intelligent cyber attack has been designed with reinforcement learning techniques to corrupt the reference high-level velocity signals in a manner to provoke accidents among the string of CAVs as soon as possible. Several real-data simulations have been performed for a string of CAVs, showing that although the cyber attack design has been successful in causing accidents without the speed analysis layer, the presented eco-cruise control could guarantee safety for the CAVs even in case of a more sophisticated cyber attack.

REFERENCES

Bangquan, X. and Xiao Xiong, W. (2019). Real-time embedded traffic sign recognition using efficient convolutional neural network. *IEEE Access*, 7, 53330–53346.

Bougiouklis, A., Korkofigkas, A., and Stamou, G. (2018). Improving fuel economy with LSTM networks and reinforcement learning. *Artificial Neural Networks and Machine Learning - ICANN 2018*, 230–239.

Dong, C., Wang, H., Ni, D., Liu, Y., and Chen, Q. (2020). Impact evaluation of cyber-attacks on traffic flow of connected and automated vehicles. *IEEE Access*, 8, 86824–86835.

Gáspár, P. and Németh, B. (2019). *Predictive Cruise Control for Road Vehicles Using Road and Traffic Information*. Springer Verlag.

Guo, L., Ye, J., and Yang, B. (2021). Cyberattack detection for electric vehicles using physics-guided machine learning. *IEEE Transactions on Transportation Electrification*, 7, 2010–2022.

Guo, Q., Li, L., and (Jeff) Ban, X. (2019). Urban traffic signal control with connected and automated vehicles: A survey. *Transportation Research Part C: Emerging Technologies*, 101, 313–334.

Hellström, E., Ivarsson, M., Åslund, J., and Nielsen, L. (2009). Look-ahead control for heavy trucks to minimize trip time and fuel consumption. *Control Engineering Practice*, 17(2), 245–254.

Jeon, W., Xie, Z., Zemouche, A., and Rajamani, R. (2021). Simultaneous cyber-attack detection and radar sensor health monitoring in connected acc vehicles. *IEEE Sensors Journal*, 21, 15741–15752.

Kim, K., Kim, J.S., Jeong, S., Park, J.H., and Kim, H.K. (2021). Cybersecurity for autonomous vehicles: Review of attacks and defense. *Computers & Security*, 103, 1–27.

Li, Y., Tu, Y., Fan, Q., Dong, C., and Wang, W. (2018). Influence of cyber-attacks on longitudinal safety of connected and automated vehicles. *Accident Analysis & Prevention*, 121, 148–156.

Liu, L., Tang, X., Xie, J., Gao, X., Zhao, W., Mo, F., and Zhang, G. (2020). Deep-learning and depth-map based approach for detection and 3-d localization of small traffic signs. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 13, 2096–2111.

Liu, T., Tian, H., Tian, G., and Huang, Y. (2017). Neural network based online eco-driving strategy for plug-in hybrid electric bus. In *2017 IEEE Vehicle Power and Propulsion Conference (VPPC)*, 1–5.

Németh, B. and Gáspár, P. (2017). The relationship between the traffic flow and the look-ahead cruise control. *IEEE Transactions on Intelligent Transportation Systems*, 18(5), 1154–1164.

Németh, B., Mihály, A., and Gáspár, P. (2021). Design of fault-tolerant cruise control in a hierarchical frameworkfor connected automated vehicles. *5th International Conference on Control and Fault-Tolerant Systems (SysTol'21)*, 1–6.

Padilla, G.P., Weiland, S., and Donkers, M.C.F. (2018). A global optimal solution to the eco-driving problem. *IEEE Control Systems Letters*, 2(4), 599–604.

Passenberg, B., Kock, P., and Stursberg, O. (2009). Combined time and fuel optimal driving of trucks based on a hybrid model. *European Control Conference, Budapest*.

Saerens, B., Rakha, H., Diehl, M., and den Bulck, E.V. (2013). A methodology for assessing eco-cruise control for passenger vehicles. *Transportation Research Part D*, 19, 20–27.

Sciarretta, A. and Vahidi, A. (2019). *Energy-Efficient Driving of Road Vehicles*. Springer Verlag.

Wu, G., Ye, F., Hao, P., Esaid, D., Boriboonsomsin, K., and Barth, M. (2019). Deep learning-based eco-driving system for battery electric vehicles.

Zhu, J., Ngo, C., and Sciarretta, A. (2019). Real-time optimal eco-driving for hybrid-electric vehicles. *IFAC-PapersOnLine*, 52(5), 562 – 567. 9th IFAC Symposium on Advances in Automotive Control AAC 2019.