Pro-social control of connected automated vehicles in mixed-autonomy multi-lane highway traffic

Downloaded from: https://research.chalmers.se, 2025-12-04 18:26 UTC

Citation for the original published paper (version of record):

Larsson, J., Keskin, F., Peng, B. et al (2021). Pro-social control of connected automated vehicles in mixed-autonomy multi-lane highway traffic. Communications in Transportation Research, 1(December). http://dx.doi.org/10.1016/j.commtr.2021.100019

N.B. When citing this work, cite the original published paper.

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library

(article starts on next page)

Communications in Transportation Research 1 (2021) 100019

Contents lists available at ScienceDirect

Communications in Transportation Research

journal homepage: www.journals.elsevier.com/communications-in-transportation-research

Pro-social control of connected automated vehicles in mixed-autonomy multi-lane highway traffic

Jacob Larsson [a], Musa Furkan Keskin [a], Bile Peng [b], Balázs Kulcsár [a,*], Henk Wymeersch [a]

[a] Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, 41 296, Sweden
[b] Institute for Communications Technology, TU Braunschweig, Braunschweig, 38 106, Germany

ARTICLE INFO

Keywords:
Altruistic control
Pro-social control
Traffic disturbance
Model predictive control
Connected automated vehicles
Stop-and-go waves

ABSTRACT

We propose pro-social control strategies for connected automated vehicles (CAVs) to mitigate jamming waves in mixed-autonomy multi-lane traffic, resulting from car-following dynamics of human-driven vehicles (HDVs). Different from existing studies, which focus mostly on ego vehicle objectives to control CAVs in an individualistic manner, we devise a pro-social control algorithm. The latter takes into account the objectives (i.e., driving comfort and traffic efficiency) of both the ego vehicle and surrounding HDVs to improve smoothness of the entire observable traffic. Under a model predictive control (MPC) framework that uses acceleration and lane change sequences of CAVs as optimization variables, the problem of individualistic, altruistic, and pro-social control is formulated as a non-convex mixed-integer nonlinear program (MINLP) and relaxed to a convex quadratic program through converting the piece-wise-linear constraints due to the optimal velocity with relative velocity (OVRV) car-following model into linear constraints by introducing slack variables. Low-fidelity simulations using the OVRV model and high-fidelity simulations using PTV VISSIM simulator show that pro-social and altruistic control can provide significant performance gains over individualistic driving in terms of efficiency and comfort on both single- and multi-lane roads.

1. Introduction

In urban transportation systems, traffic jams pose a significant threat to vehicle safety, exhaust gas emission, fuel economy and passenger comfort, especially in dense traffic scenarios with stop-and-go waves. Disturbances such as accidents, lane restrictions or random braking may propagate backwards through the traffic as a result of the car-following dynamics of human-driven vehicles (HDVs), which leads to moving traffic jams Sugiyama et al. (2008); Wang et al. (2016); Stern et al. (2018); Dabiri and Kulcsár (2017); Pereira et al. (2022). To reduce such disturbance propagation, connected automated vehicles (CAVs) can be applied to control traffic flow and smooth out stop-and-go waves Kamal et al. (2014, 2016); Wang et al. (2016); Stern et al. (2018); Luspay et al. (2010). To circumvent jamming waves on single-lane roads, predictive control of CAV acceleration has recently been a popular strategy Kamal et al. (2014); Dollar and Vahidi (2018); Stern et al. (2018); Csikos and Kulcsar (2017); Dabiri and Kulcsár (2017); Luspay et al. (2012), while, for multi-lane highways, high-level lane changing decisions can be incorporated as additional degrees of freedom that can be optimized jointly with low-level acceleration inputs Yu et al. (2019).

Historically in the literature, the emphasis of autonomous driving has been focused on CAVs' own (selfish) driving objectives Kamal et al. (2016); Bahram (2017); Yu et al. (2019), while ignoring the traffic-smoothing properties of CAVs Stern et al. (2018). In Kamal et al. (2016), a model predictive control (MPC) framework is proposed, where efficiency, comfort and safety of the CAV is improved by optimizing the acceleration and lane changes in a multi-lane traffic scenario. However, the focus of Kamal et al. (2016) is solely on the driving objectives of the CAV and not the surroundings, making such a driving strategy *selfish*. Another MPC-based approach is presented in Bahram (2017), where a mixed-integer quadratic programming problem is setup to optimize longitudinal velocity and lane-change maneuvers of the CAV. An entirely different selfish control strategy using reinforcement learning (RL) is presented in Yu et al. (2019), where a multi-agent (multi-vehicle) RL algorithm is trained to achieve coordination between multiple CAVs in a highway scenario.

Altruistic agents have been considered in a variety of fields, including social dynamics analysis via linear quadratic Gaussian (LQG) control Huang et al. (2010), traffic route management Levy and Ben-Elia (2016); Bıyık et al. (2018), microscopic traffic control Wang et al. (2017), water

* Corresponding author.
E-mail address: kulcsar@chalmers.se (B. Kulcsár).

https://doi.org/10.1016/j.commtr.2021.100019
Received 29 September 2021; Received in revised form 18 November 2021; Accepted 18 November 2021
Available online 30 November 2021
2772-4247/© 2021 The Author(s). Published by Elsevier Ltd on behalf of Tsinghua University Press. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

resource planning via Markov decision processes (MDP) Desquesnes et al. (2017) and uncertain dynamic games Camponogara et al. (2006). In Levy and Ben-Elia (2016), a macroscopic routing perspective is presented to compare the total driving time of vehicles in a network obtained by a *selfish* user equilibrium (UE) model and an *altruistic* social optimum (SO) model. Similarly, the work in Bıyık et al. (2018) provides a game-theoretic analysis of altruistic autonomy from a vehicle routing perspective and investigates its effect on traffic latency under varying degrees of altruism of CAVs. Regarding microscopic control, a cooperative altruistic driving strategy is developed in Wang et al. (2017), where traffic jamming on highways is resolved by coordinating a group of CAVs using vehicle-to-vehicle (V2V) communications.

In light of the existing literature on traffic control, implementing purely individualistic (selfish) and/or purely altruistic behaviour have partially been investigated before. While a vehicle that behaves selfish/individualistic allocates control input to actions to reward only itself, pure altruistic behaviour allocates input in a way to reward only others (and disregard its own rewards or benefits). These forms of expected CAV behaviours can often be conflicting, even though they may be beneficial in different traffic scenarios. Therefore, one possible way to overcome the behaviour dilemma (selfish or altruistic?) is to use both at the same time. As shown in Fig. 1, the key insight of the altruistic controller is that performing altruistic lane change maneuvers help to dissolve jamming waves while improving comfort and efficiency. In Buckman et al. (2019); Pierson et al. (2020), socially compliant central coordination algorithms are suggested to solve a traffic coordination problem. In particular, these studies propose a large variety of Social Value Orientation (SVO) algorithms via the definition of two independent reward metrics: reward to self and reward to others. In case of intersection crossing, in Buckman et al. (2019), pro-social (combination of selfish and altruistic) behaviour provides most of the benefits in terms of wait time reduction. Mavrogiannis et al. (2020) proposes a decentralized intersection coordination mechanism using principles similar to Buckman et al. (2019): topological braids capture selfish-altruistic modes. In the latter, however, elimination of unsafe trajectories is the goal. SVOs are mapped via proper weighting strategies, emphasizing the relative importance of altruistic or selfish objectives.

The aim with the above state-of-the-art methods to influence/control/coordinate automated vehicles is the same: adapt their behaviour in mixed traffic conditions. As in mixed-autonomy traffic scenarios, both HDVs and CAVs have to co-exist and CAVs may be coordinated following social values known for human drivers.

In this paper, extending our preliminary work in Keskin et al. (2020) with comprehensive low- and high-fidelity simulation results, we propose in a comparative environment *altruistic* and *pro-social* control/coordination strategies where CAVs mitigate traffic jams by optimizing the driving objectives of the overall traffic as well as of their own selfish trajectories. We define the optimization problem that empowers CAVs with social behaviours as a model based, finite horizon, multi-objective optimization problem.

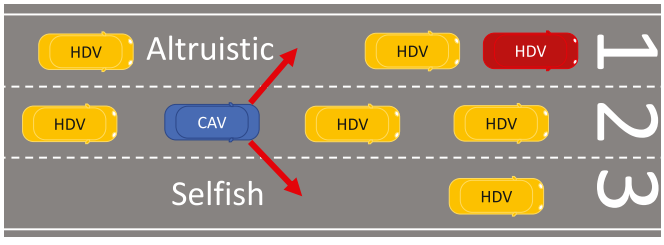

Fig. 1. Exemplary multi-lane highway scenario with CAVs (blue) and HDVs (yellow) and decelerating HDVs (red). An altruistic strategy would involve turning left in to lane 1 and try to mitigate the traffic jamming caused by the red HDV by acting as a damper and help improve overall traffic smoothness. A selfish decision is to take a right on to lane 3 and avoid the jamming, which would improve smoothness for the CAV itself, but lead to jamming on lane 1. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

Therefore, the contribution of this paper is twofold:

- We develop solutions to incentivize the pro-social coordination of mixed autonomy vehicles with model based predictive optimization algorithms. To that end, we use proper weighting strategies to map SVOs into vehicle control solutions (traffic efficiency and comfort). The key idea is to emphasize the relative importance of altruistic versus selfish objectives and reach pro-social behavior.
- We evaluate the proposed methods on both low- and high-fidelity traffic simulators, which helps us quantify the benefits/drawbacks of SVOs in terms of fuel economy, ride comfort, trajectory alignment, etc.

More precisely, in this work, we propose an MPC-based selfish, altruistic, and pro-social coordination algorithm, where CAVs in mixed traffic scenarios model HDVs with the Optimal Velocity with Relative Velocity (OVRV) car-following model Wilson and Ward (2011). Then, a finite horizon time prediction window is created in which CAVs can select their speed and lanes to minimize the multi-objective cost function. The result is a non-convex mixed-integer nonlinear program, which is relaxed to a simpler convex quadratic program via penalty based reformulation of the OVRV model. Simulations with low fidelity simulations by the OVRV model, and with high fidelity simulations by PTV VISSIM microscopic traffic simulator are carried out to investigate the impact of different social behaviour triggered by CAVs.

2. System model

In Fig. 1, a mixed autonomy multi-highway traffic scenario is depicted, with both CAVs and HDVs. It is assumed that the CAVs obtain position and speed information from surrounding HDVs[1] via vehicle-to-vehicle (V2V) communications Zhang and Orosz (2018); Avedisov et al. (2020). Although there may be any number of CAVs in the traffic, the focus here lies on individual automated driving, i.e., CAVs take decision without cooperating explicitly with other CAVs. In the scenario, the objective of individual CAVs is to obtain optimal control input sequences in terms of vehicle acceleration input and lane change decisions. The CAV applies an altruistic MPC controller and thus the optimization objective is to maximize the entire traffic objectives in terms of comfort, efficiency and indirectly, emissions.

2.1. Vehicle states

The state vector of ith CAV at discrete time k, with sampling time Δt, is defined as

$$\mathbf{x}_{i,k}^{\mathrm{CAV}} = \left[p_{i,k}^{\mathrm{CAV}} \; v_{i,k}^{\mathrm{CAV}} \; y_{i,k}^{\mathrm{CAV}}\right]^{\mathrm{T}} \tag{1}$$

for $i = 1, \ldots, N_{\mathrm{CAV}}$, where $p_{i,k}^{\mathrm{CAV}} \in \mathbb{R}$ and $v_{i,k}^{\mathrm{CAV}} \in \mathbb{R}$ are, respectively, the longitudinal position and velocity of the vehicle, and $y_{i,k}^{\mathrm{CAV}} \in \mathcal{L} \triangleq \{1, 2, \ldots, N_{\mathrm{lane}}\}$ represents the lane number of the vehicle. Similarly, the state vector of the jth HDV at time k is expressed as

$$\mathbf{x}_{j,k}^{\mathrm{HDV}} = \left[p_{j,k}^{\mathrm{HDV}} \; v_{j,k}^{\mathrm{HDV}} \; y_{j,k}^{\mathrm{HDV}}\right]^{\mathrm{T}} \tag{2}$$

for $j = 1, \ldots, N_{\mathrm{HDV}}$, where $p_{j,k}^{\mathrm{HDV}}, v_{j,k}^{\mathrm{HDV}} \in \mathbb{R}$ and $y_{j,k}^{\mathrm{HDV}} \in \mathcal{L}$.

[1] The set of surrounding HDVs for a given CAV consists of HDVs that reside in the communication range of the CAV. Hence, HDVs with no line-of-sight (LOS) link to the CAV can also belong to this set if V2V communication between the two vehicles is possible.

2.2. CAV control inputs

The control input vector of the ith CAV at time k is given by

$$\mathbf{u}_{i,k}^{\mathrm{CAV}} = \left[a_{i,k}^{\mathrm{CAV}} \ \delta_{i,k}^{\mathrm{CAV}}\right]^{\mathrm{T}} \tag{3}$$

where $a_{i,k}^{\mathrm{CAV}} \in \mathbb{R}$ is the longitudinal acceleration and $\delta_{i,k}^{\mathrm{CAV}}$ represents the lateral movement, i.e., the lane change decision, defined as

$$\delta_{i,k}^{\mathrm{CAV}} \in \mathcal{L}_{\Delta} \triangleq \{-1, 0, 1\} \tag{4}$$

where 0 denotes a lane-keeping decision and $1/-1$ represents a left/right lane change decision. To reduce modelling complexity in controller design step (prediction), it is assumed that the lane change is instantaneous and is thus completed in a single time step Kheterpal et al. (2018). Note, more complex and dynamic lane change maneuvers can be added if needed. Finally, in the numerical case study, we use VISSIM that has a continuous and dynamic lane change model.

2.3. Car-following behavior of HDVs

The longitudinal dynamics of HDVs is described using a car-following model as in Orosz et al. (2010, 2009)

$$a_{j,k}^{\mathrm{HDV}} = f(h_{j,k}^{\mathrm{HDV}}, v_{j,k}^{\mathrm{HDV}}, \Delta v_{j,k}^{\mathrm{HDV}}) \tag{5}$$

where $a_{j,k}^{\mathrm{HDV}} \in \mathbb{R}$ is the longitudinal acceleration of the jth HDV at time k, $h_{j,k}^{\mathrm{HDV}}$ is the headway and $\Delta v_{j,k}^{\mathrm{HDV}}$ the velocity difference between the jth HDV and the preceding vehicle, written as

$$h_{j,k}^{\mathrm{HDV}} = p_{j,k}^{\mathrm{HDV,pre}} - p_{j,k}^{\mathrm{HDV}} \ , \tag{6}$$

$$\Delta v_{j,k}^{\mathrm{HDV}} = v_{j,k}^{\mathrm{HDV,pre}} - v_{j,k}^{\mathrm{HDV}} \ , \tag{7}$$

with $p_{j,k}^{\mathrm{HDV,pre}}$ and $v_{j,k}^{\mathrm{HDV,pre}}$ representing the position and speed of the vehicle preceding the jth HDV at time k on the same lane. The car-following dynamics are represented by the Optimal Velocity with Relative Velocity (OVRV) model Wilson and Ward (2011)

$$f(h, v, \Delta v) = \alpha(V(h) - v) + \beta \Delta v \tag{8}$$

where the velocity function $V(h)$ is a piecewise-linear function of headway h (driver perceived optimal and headway based velocity), defined as Zhang and Orosz (2016)

$$V(h) = \left[\tilde{V}(h)\right]_0^{v_{\max}}, \tilde{V}(h) = v_{\max}\frac{h - h_{\min}}{h_{max} - h_{\min}} \ , \tag{9}$$

with $[v]_0^{v_{\max}} \triangleq \max(0, \min(v_{\max}, v))$, and α, β, $h_{\min}$, $h_{\max}$ and $v_{\max}$ are driver-dependent model parameters. Furthermore, it is also assumed that HDVs keep the same lane over the entire MPC prediction horizon Kamal et al. (2016), i.e., $\delta_{j,k+n}^{\mathrm{HDV}} = 0$ for $n = 0, 1, \ldots, N_{\mathrm{p}} - 1$, where N_{p} denotes the prediction horizon. In Fig. 1, the described car-following behaviour is shared among the yellow HDVs. However, the red HDV causing disturbance on lane 1 is considered a leading HDV on that lane. These HDVs are handled differently, where we implement MPC prediction heuristics in the form of constant acceleration over the horizon.

2.4. Discrete-time vehicle dynamics

The dynamics of the ith CAV can be expressed as

$$\mathbf{x}_{i,k+1}^{\mathrm{CAV}} = \mathbf{A}\mathbf{x}_{i,k}^{\mathrm{CAV}} + \mathbf{B}\mathbf{u}_{i,k}^{\mathrm{CAV}} \tag{10}$$

where

$$\mathbf{A} = \begin{bmatrix} 1 & \Delta t & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \ \mathbf{B} = \begin{bmatrix} \Delta t^2/2 & 0 \\ \Delta t & 0 \\ 0 & 1 \end{bmatrix} \tag{11}$$

In a similar fashion, the dynamics of the jth HDV can be written as

$$\mathbf{x}_{j,k+1}^{\mathrm{HDV}} = \mathbf{A}\mathbf{x}_{j,k}^{\mathrm{HDV}} + \mathbf{B}\mathbf{u}_{j,k}^{\mathrm{HDV}} \tag{12}$$

where the input is defined as

$$\mathbf{u}_{j,k}^{\mathrm{HDV}} = \left[a_{j,k}^{\mathrm{HDV}} \ \delta_{j,k}^{\mathrm{HDV}}\right]^{\mathrm{T}} . \tag{13}$$

Remark 1. By assumption, HDVs are either equipped with communication devices (e.g., app, specific hardware, V2I or V2V communication device) or the infrastructure can sense HDVs (e.g. camera, 5G). In both of the cases, we consider HDVs to actively or passively share with CAVs some of their dynamic variables (position, speed).

Remark 2. CAVs require to have access to position and speed variables of the surrounding vehicles. Pro-social receding control algorithm can then be run by a CAV (e.g., on-board optimizer units). This means that CAVs do not necessary need to share any further data among themselves. Note, that cooperative pro-social control algorithms create future research directions (see Conclusions).

3. MPC formulation for individual altruistic driving

This section covers the problem of how altruism can be reached with carefully selecting CAV control input. The constraints for inputs and states are provided for a multi-lane traffic scenario, and then the optimal CAV control problem is formulated in the MPC framework.

3.1. Constraints

For the optimal CAV control problem, the following constraints are imposed on the vehicle inputs and states.

3.1.1. Acceleration bounds

The following constraints bound the longitudinal acceleration, i.e.,

$$a_{\min} \leq a_{i,k+n}^{\mathrm{CAV}} \leq a_{\max} \ , \ n = 0, 1, \ldots, N_{\mathrm{p}} - 1 \tag{14}$$

where N_{p} is the horizon length.

3.1.2. Lateral safety constraints

At the nth prediction step, a lane change occurs when $\delta_{i,k+n}^{\mathrm{CAV}} = 1$. Here, the ith CAV should keep the headway h_{safe} to the closest vehicle in the new lane, i.e.,

$$p_{i,k+n}^{\mathrm{CAV,bg}} - p_{i,k+n}^{\mathrm{CAV}} \geq h_{\mathrm{safe}}, \ p_{i,k+n}^{\mathrm{CAV}} - p_{i,k+n}^{\mathrm{CAV,sm}} \geq h_{\mathrm{safe}} \tag{15}$$

where $p_{i,k+n}^{\mathrm{CAV,bg}}$ and $p_{i,k+n}^{\mathrm{CAV,sm}}$ are the longitudinal positions of the vehicles on the new lane that are closest to the ith CAV at time $k + n$ with $p_{i,k+n}^{\mathrm{CAV,bg}} \geq p_{i,k+n}^{\mathrm{CAV}} \geq p_{i,k+n}^{\mathrm{CAV,sm}}$. Furthermore, we also limit the number of lane changes over the horizon to at most 1.

$$\sum_{n=1}^{N_p - 1} \delta_{i,k+n}^{\mathrm{CAV}} \leq 1 \tag{16}$$

3.1.3. Longitudinal safety constraints

In order to avoid collisions and keep a safe minimum headway to preceding vehicles on the lane, CAVs and HDVs are constrained to a dynamic headway Kamal et al. (2014). Hence, for the ith CAV, we have

$$p_{i,k+n}^{\mathrm{CAV,pre}} - p_{i,k+n}^{\mathrm{CAV}} \geq h_{\min} + t_{\min} v_{i,k+n}^{\mathrm{CAV}} \ , \ n = 1, \ldots, N_{\mathrm{p}} \tag{17}$$

where $p_{i,k+n}^{\mathrm{CAV,pre}}$ is the position of the vehicle preceding the ith CAV at time

$k+n$ and $t_{\min}$ denotes the minimum time headway. Similar constraints also bound all HDVs during the prediction in order to augment the OVRV model behaviour with a dynamic and realistic headway behaviour.

3.2. Objectives

For the optimal CAV control problem, we apply two categories of objectives, namely, traffic efficiency and driving comfort. In addition, we split the driving comfort into acceleration and jerk components.

3.2.1. Traffic efficiency

Traffic efficiency can be defined as the objective of maintaining a desired velocity V^* for the ith CAV and for the overall traffic including observable HDVs and other CAVs:

$$\begin{aligned}&\mathcal{J}_i^{\text{eff}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})\\&=\sum_{n=0}^{N_p-1}\left[\left(v_{i,k+n}^{\text{CAV}}-V^*\right)^2+\kappa\sum_{j\in\mathcal{G}_{i,k}^{\text{HDV}}}\left(v_{j,k+n}^{\text{HDV}}-V^*\right)^2\right]\end{aligned}\tag{18}$$

where $\mathcal{G}_{i,k}^{\text{HDV}}$ is the set of indices of those HDVs succeeding the ith CAV at time k and that are also observable by it,[2] and κ is a constant variable that indicates the *level of altruism*, i.e., a weight that controls the CAV's prioritization between its own selfish driving objectives and the surrounding traffics objectives.[3] Finally, V^* stands for the expected target speed (lane-wise speed limit or equilibrium speed). The dependency of HDV velocities on CAV control inputs is through Eqs. (5)–(7).

3.2.2. Driving comfort - acceleration magnitude

This driving objective aims to reduce the discomfort associated with large magnitudes of acceleration:

$$\begin{aligned}&\mathcal{J}_i^{\text{mag}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})\\&=\sum_{n=0}^{N_p-1}\left[\left(a_{i,k+n}^{\text{CAV}}\right)^2+\kappa\sum_{j\in\mathcal{G}_{i,k}^{\text{HDV}}}\left(a_{j,k+n}^{\text{HDV}}\right)^2\right]\end{aligned}\tag{19}$$

where the dependency of HDV accelerations on CAV control inputs is through Eqs. (5)–(7).

3.2.3. Driving comfort - jerk

Another component of discomfort associated with acceleration is the jerky behaviour due to rapid changes in the acceleration derivative:

$$\begin{aligned}\mathcal{J}_i^{\text{jerk}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})=\sum_{n=0}^{N_p-1}&\left[\left(\frac{a_{i,k+n+1}^{\text{CAV}}-a_{i,k+n}^{\text{CAV}}}{\Delta t}\right)^2\right.\\&\left.+\kappa\sum_{j\in\mathcal{G}_{i,k}^{\text{HDV}}}\left(\frac{a_{j,k+n+1}^{\text{HDV}}-a_{j,k+n}^{\text{HDV}}}{\Delta t}\right)^2\right]\end{aligned}\tag{20}$$

which is an approximation of the functions derivative obtained by Euler's method.

3.2.4. Driving comfort - total objective

We define the total driving comfort as a two-component objective, where the individual parts above are weighted against each other. The total driving comfort objective for the ith CAV at time k is defined as

$$\begin{aligned}&\mathcal{J}_i^{\text{comf}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})\\&=\mathcal{J}_i^{\text{mag}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})+w_2\,\mathcal{J}_i^{\text{jerk}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})\end{aligned}\tag{21}$$

where w_2 is a weight to indicate the importance of the jerk component on to the overall comfort objective.

3.2.5. Total objective function

For the ith CAV, the total objective function at time k with prediction horizon length N_p is defined as

$$\begin{aligned}&\mathcal{J}_i^{\text{tot}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})\\&=\mathcal{J}_i^{\text{eff}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})+w_1\,\mathcal{J}_i^{\text{comf}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})\end{aligned}\tag{22}$$

where w_1 is a weight that balances the impact between efficiency and comfort objectives. We note that safety is taken into consideration as hard physical constraints through Eq. (15) and Eq. (17) in the MPC formulation.

3.3. MPC prediction heuristics

Over the MPC prediction horizon, we assume a constant acceleration heuristic Hamdar et al. (2008); Kesting et al. (2010); M Treiber (2013); Yu et al. (2021) for predicting the leading HDV trajectory.[4] At time k, we have

$$a_{j,k+n}^{\text{HDV}}=\hat{a}_{j,k}^{\text{HDV}},\ j\in\mathcal{F}_{i,k}^{\text{HDV}}\tag{23}$$

for $n=0,1,\ldots,N_p-1$, where $\mathcal{F}_{i,k}^{\text{HDV}}$ is the set of indices for leading HDVs that are observed by the ith CAV, and $\hat{a}_{j,k}^{\text{HDV}}$ is the measured acceleration of the jth leading HDV at time k. To ensure the non-negativity of the velocity over the prediction horizon in Eq. (23), $a_{j,k+n}^{\text{HDV}}$ is set to zero for $n>\tilde{n}$ if $v_{j,k+\tilde{n}}^{\text{HDV}}<0$. In other words, we prioritize constraining the velocity over Eq. (23) in the prediction horizon.

3.4. Problem formulation

Given the initial internal states of the ith CAV at time k and the initial states of HDVs observed by the ith CAV at time k, the MPC problem over a prediction horizon of length N_p can be formulated as follows:

$$\begin{aligned}&\underset{\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}}}{\text{minimize}}\ \mathcal{J}_i^{\text{tot}}(\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}})\\&\text{subjectto}\\&(\text{Prediction of Leading HDVs})\,Eq.\ (23)\\&(\text{HDV Car}-\text{Following Model})\,Eqs.\ (5)\text{–}(9)\\&(\text{CAV}-\text{HDV Dynamics})\,Eqs.\ (10)\text{–}(13)\\&(\text{Vehicle/TrafficConstraints})\,Eqs.\ (14)\text{–}(17)\ .\end{aligned}\tag{24}$$

The formulation in Eq. (24) is a mixed-integer non-linear programming problem (MINLP). The non-linearity (and non-convexity) stem from the car-following dynamics of the HDVs in Eq. (8), where the range policy Eq. (9) is a piecewise-linear function. The integer variables come from the lane change sequence variables $\delta_{i,k:k+N_p-1}^{\text{CAV}}$.

[2] The CAV is only able to control succeeding vehicles, i.e., Lagrangian control scheme Stern et al. (2018)), utilizing state information from both preceding and succeeding vehicles.

[3] We note that both CAV and HDV velocities in Eq. (18) depend on CAV control inputs $\mathbf{u}_{i,k:k+N_p-1}^{\text{CAV}}$ through Eqs. (5)–(7) and Eq. (10). From the car-following behavior in Eqs. (5)–(7), HDV acceleration is a function of the position and speed of the preceding vehicle, which implies that the effect of CAV control actions can be propagated downstream towards HDVs moving on the same lane and affect the traffic efficiency in Eq. (18).

[4] Accelerations of leading HDVs on each lane cannot be determined using a car-following function as in Eq. (5). Thus, we assume the accelerations are available through on-board tracking filters on CAVs and therefore the accelerations can be used to predict leading HDVs trajectories.

4. Optimization strategies for altruistic driving

In this section, we propose an optimization strategy for handling the MINLP in Eq. (24). This is achieved by firstly reformulating the non-linear car-following dynamics in Eq. (8) as linear constraints. To handle the integer lane change variables in $\delta^{\text{CAV}}_{i,k:k+N_\text{p}-1}$, we decompose the optimization problem into lower-level subproblems on each reachable lane. The solutions of these problems can then be combined to form a final lane-change decision.

4.1. Transformation of piecewise-linear car-following constraints

To avoid the intractability of the piecewise-linear car-following dynamics in Eqs. (5)–(9), we take a penalty based approach for constraint reformulation. Precisely, we transform the piecewise-linear equality constraint in Eq. (5) to a linear equality constraint, two linear inequality constraints that bound the function from above and below, along with a penalty term in the objective function. The car-following dynamics constraint in Eq. (5) with the OVRV model in Eq. (8), given by

$$a^{\text{HDV}}_{j,k+n} = \alpha(V(h^{\text{HDV}}_{j,k+n}) - v^{\text{HDV}}_{j,k+n}) + \beta\Delta v^{\text{HDV}}_{j,k+n} \tag{25}$$

for $n = 0, 1, \ldots, N_\text{p} - 1$ and $j \in \mathcal{H}^{\text{HDV}}_{i,k}$, where $\mathcal{H}^{\text{HDV}}_{i,k}$ is a set of indices for HDVs observed by the ith CAV, except for the leading HDVs on each lane, can be rewritten by introducing slack variables $\boldsymbol{\gamma}_j = \left[\gamma_{j,N_\text{p}-1}, \ldots, \gamma_{j,0}\right]^\text{T}$ as follows. We begin by introducing a penalty term in the objective function Eq. (24) using the ℓ_2-norm as

$$\mathcal{J}^{\text{tot}}_i(\mathbf{u}^{\text{CAV}}_{i,k:k+N_\text{p}-1}) + \lambda\,\mathcal{J}^{\text{slack}} \tag{26}$$

where

$$\mathcal{J}^{\text{slack}} = \sum_{j\in\mathcal{H}^{\text{HDV}}_{i,k}} |\boldsymbol{\gamma}_j|^2 \,. \tag{27}$$

and $\boldsymbol{\gamma}_j = \left[\gamma_{j,0} \cdots \gamma_{j,N_\text{p}-1}\right]^\text{T}$. Here, λ is a weight that controls the tightness of car-following dynamics in Eq. (25) and $\mathcal{J}^{\text{slack}}$ enforces the driving models' slack variable to be minimized. Secondly, we reformulate Eq. (25) by introducing one equality constraint and two inequality constraints that bound the function from above and below, as

$$a^{\text{HDV}}_{j,k+n} = \alpha(\tilde{V}(h^{\text{HDV}}_{j,k+n}) - v^{\text{HDV}}_{j,k+n}) + \beta\Delta v^{\text{HDV}}_{j,k+n} + \gamma_{j,n} \tag{28}$$

$$a^{\text{HDV}}_{j,k+n} \le \alpha(\tilde{V}(h_{\max}) - v^{\text{HDV}}_{j,k+n}) + \beta\Delta v^{\text{HDV}}_{j,k+n} \tag{29}$$

$$a^{\text{HDV}}_{j,k+n} \ge \alpha(\tilde{V}(h_{\min}) - v^{\text{HDV}}_{j,k+n}) + \beta\Delta v^{\text{HDV}}_{j,k+n} \tag{30}$$

for $n = 0, 1, \ldots, N_\text{p} - 1$ and $j \in \mathcal{H}^{\text{HDV}}_{i,k}$.

Remark 3. Equation 28, 29 and 30 are linear constraints that we use in order to replace the non-linear function from equation from equation 5 – 9. This is an approximation of the true driving model (OVRV), but by using a penalty function (equation 27) and minimize the slack variable in equation 28, we in fact can get very close to the OVRV driving model.

4.2. Optimization subproblem for fixed lane change decision

In order to handle the integer lane change variables in Eq. (24), we create three separate subproblems for each reachable lane for the ith CAV, where each subproblem corresponds to a fixed lane changing decision $\delta^{\text{CAV}}_{i,k} \in \mathcal{L}_\Delta$. This means that each lane change decision is optimized only for the initial time of the MPC control problem, while the acceleration control inputs are still obtained for N_p steps forward in time on each lane. With this reformulation in Eqs. 26–30, the MPC optimization subproblem for each reachable lane of Eq. (24) for a given lane change decision can be written as:

$$\begin{aligned}
&\underset{a^{\text{CAV}}_{i,k:k+N_\text{p}-1},\{\gamma_j\}}{\text{minimize}}\ \mathcal{J}^{\text{tot}}_i(\mathbf{u}^{\text{CAV}}_{i,k:k+N_\text{p}-1}) + \lambda\,\mathcal{J}^{\text{slack}}\\
&\text{subjectto}\\
&(\text{Prediction of Leading HDVs})\textit{Eq.}\ (23)\\
&(\text{HDV Car}-\text{Following Model})\textit{Eqs.}\ (28)\text{–}(30)\\
&(\text{CAV}-\text{HDV Dynamics})\textit{Eqs.}\ (10)\text{–}(13)\\
&(\text{Vehicle/TrafficConstraints})\textit{Eqs.}\ (14)\text{–}(17)\ .
\end{aligned} \tag{31}$$

Note that Eq. (31) is a convex optimization problem, with convex quadratic objective function and linear constraints. Thus, the optimization problem can be solved efficiently using interior-point methods Boyd and Vandenberghe (2004). The solution of Eq. (31) for the three different lane sub-problems can then be compared against each other, and the lane change decision with the lowest cost with its corresponding acceleration sequence is chosen as the optimal control inputs.

Remark 4. With the above relaxations complexity of the modeling and optimization are kept low for two main reasons. First, the dynamism of the lane change is faster in our setup than the longitudinal changes of variables. Second, the the runtime of the control algorithm is acceptable.

Remark 5. No lane change maneuvers are executed by the HDVs in the prediction model. CAVs can decide the lane they take at the begining of the optimization. Note, however, in the case of an HDV changing lane, the prediction model is updated with the new lane position. Since the controller uses the receding horizon paradigms, this update is done at every sampling time.

5. Experiments

We perform experiments at two levels of simulation fidelity. First, we carry out low-fidelity simulations where the simpler OVRV model in Eq. (8) is used to simulate the car-following dynamics of HDVs in evaluating the performance of the proposed MPC controller in Eq. (31) on a single lane scenario. Then, the high-fidelity microscopic multi-modal traffic simulator PTV VISSIM is deployed to verify the controller in realistic single- and multi-lane settings. Whilst with the OVRV model we validate the proof of concept in simplified traffic scenarios, with the multi-lane scenario VISSIM tools, we partially test the nominal controller in uncertain (car-following model mismatch) and more complex and realistic traffic environments.

5.1. Simulation setup and parameters

We consider three simulation scenarios with two different setups. Firstly, the OVRV and VISSIM single lane scenarios use a setup where there is a leading HDV driving with a sinusoidal acceleration profile.[5] Following this leading vehicle there is directly a CAV that tries to mitigate these disturbances.[6] HDV driving with the OVRV model or the VISSIM Wiedemann99 model is then following this CAV, and these HDVs are being controlled by the CAV.

For the third scenario, which is multi-lane VISSIM, there is a second setup. The road consists of three lanes and the first row of vehicles contain HDVs following a sinusoidal acceleration profile on all three lanes. Following these leading HDVs, there is the second row of another three HDVs, behaving according to the OVRV or W99 driving model, on all three lanes. Succeeding these HDVs is a single CAV in the middle lane. After the CAV comes several rows of HDVs depending on the penetration

[5] The acceleration profile reflects hard braking and acceleration with an amplitude of 5 m/s^2 with a period of 20 seconds. The HDV speed will never reach 0 m/s.

[6] The CAV is has no prior knowledge of the leading HDVs acceleration profile.

rate, e.g. 10 rows of HDVs per CAV with a penetration rate of 10%.

For all setups, HDVs that are behind the leading row of HDVs are using the OVRV or W99 model and the traffic scenarios are similar to stop-and-go traffic. The simulation is on a continuous loop track with 1 or 3 lanes for the VISSIM simulations and the OVRV model only considers a straight infinitely long road. The simulation time for the OVRV and VISSIM experiments is 200 and 800 seconds respectively. The initial headway distance for all vehicles is 40 meters for every scenario and driving model.

Both the low-fidelity OVRV and high-fidelity VISSIM simulations have several parameters that determine the behaviour of the MPC controller and HDV trajectories. In the appendix, we briefly cover the parameters pertaining to the driver model and the MPC optimization, used for the different simulation scenarios. See Appendix A for specific optimization parameters and Appendix B for specific simulation parameters.

Remark 6. The pro-social control actions can be calculated centrally, assuming high performance computers. This can be relaxed and computation can be distributed. The computational time for the rolling horizon optimization depends on the number of HDVs.

5.2. Evaluation metrics

For assessing the performance of the controller and the impact of the different levels of altruism (parameter κ), we evaluate three areas: overall traffic acceleration and velocity directly impacted by the objective functions along with vehicle emissions, which is an indirect byproduct of acceleration and velocity via the Virginia Tech microscopic (VT-Micro) vehicle emission model.

5.2.1. Vehicle emission model

Since vehicle emissions are not directly optimized in the MPC controller, we use a vehicle emission model to calculate the emissions of carbon monoxide (CO), hydrocarbons (HC), nitrogen oxide (NOx) and fuel consumption (FC). We adopt the VT-Micro model Zegeye et al. (2013) to calculate these metrics from the acceleration and velocity data gathered from simulation. In addition, it is assumed that the traffic consists solely of gasoline passenger vehicles. Under this setting, emissions can be calculated as

$$\mathcal{J}_k^y(\mathbf{v}_{i,k}, \mathbf{a}_{i,k}) = \exp(\mathbf{v}_{i,k}^{\mathrm{T}} \mathbf{P}_y \mathbf{a}_{i,k}) \tag{32}$$

where $\mathcal{J}_k^y(\mathbf{v}_{i,k}, \mathbf{a}_{i,k})$ is the prediction of the variable $y \in \{\mathrm{CO}, \mathrm{HC}, \mathrm{NOx}, \mathrm{FC}\}$ at every simulation step k of the i-th vehicle (CAV or HDV), $\mathbf{P}_y \in \mathbb{R}^{4\times 4}$ is a parameter matrix for each variable y,

$$\mathbf{v}_{i,k} = \begin{bmatrix} 1 & v_{i,k}^\star & (v_{i,k}^\star)^2 & (v_{i,k}^\star)^3 \end{bmatrix}^{\mathrm{T}} \tag{33}$$

$$\mathbf{a}_{i,k} = \begin{bmatrix} 1 & a_{i,k}^\star & (a_{i,k}^\star)^2 & (a_{i,k}^\star)^3 \end{bmatrix}^{\mathrm{T}} \tag{34}$$

are the velocity and acceleration vectors for the i-th vehicle at time k, respectively, and $\star \in \{\mathrm{CAV}, \mathrm{HDV}\}$. Emission and fuel consumption rates are given in the units kg/s and l/s, respectively.

5.3. Results

In this section, we present the experimental results for the OVRV model and the VISSIM W99 model in single- and multi-lane scenarios.

5.3.1. Single lane - OVRV

To evaluate the performance of the proposed altruistic control strategy in a single-lane road, Fig. 2 showcases the cumulative root-mean-squared (RMS) acceleration obtained via the OVRV model and the proposed approach with different values of the altruism parameter $\kappa = [0, 0.5, 1]$. It is observed that cumulative acceleration decreases with an increasing degree of altruism. The selfish controller improves upon the OVRV driving model by reducing cumulative accelerations by 3.4% while the altruistic controller with $\kappa = 1$ further improves on the results with an additional 2.1% compared to the selfish case with $\kappa = 0$. This indicates the potential of altruistic control to mitigate traffic disturbances and improve driving comfort and safety.

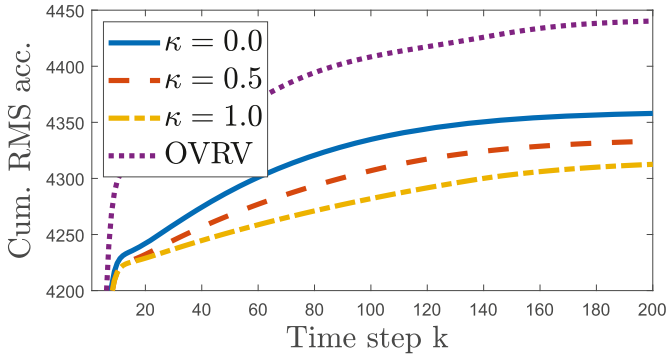

Fig. 2. Single lane OVRV simulation results of cumulative RMS acceleration with varying altruism parameter κ.

As an illustration of how stop-and-go waves are dampened via the proposed altruistic strategy, Fig. 3a shows with respect to time the headway of the HDV that follows the CAV. Comparing the OVRV driving model with the selfish case, $\kappa = 0$, there is only a minuscule smoothing effect that is barely noticeable. However, when the altruistic controller is applied, with $\kappa = 1$, the smoothing effect is more pronounced, i.e., headway fluctuations are significantly reduced, which proves the effectiveness of the altruistic strategy against the disturbances caused by the leading HDV. In Fig. 3b, we investigate the velocity fluctuations, i.e., the velocity difference between the HDV that follows the CAV and the CAV itself. It is seen that the amplitude of fluctuations decreases for the altruistic case in comparison to the OVRV driving model and the selfish controller, which again demonstrates the smoothing effect provided by the altruistic CAV. These results evidence that an altruistic CAV can reduce the oscillations stemming from the leading vehicle, leading to more stable velocities (i.e., efficiency) and lower accelerations (i.e., comfort) experienced by the vehicles.

Fig. 4a–d illustrate the benefits of altruism on the cumulative vehicle emissions, where altruism reduces the total emissions of all metrics. Similarly, Fig. 4e–h demonstrate the emission rate as distributions for the entire simulation, where the pure altruistic driving strategy exhibits a lower mean value and variance in comparison to other altruism levels.

5.3.2. Single lane - W99

In the previous subsection, no car-following model uncertainties have been considered, i.e., the MPC used the OVRV to create predictions. In PTV VISSIM, the car-following model W99 aims at capturing significantly more complex driver behaviour than the OVRV. We therefore tested the MPC algorithm (based on OVRV predictions) in a single lane context by emulating the real environment with W99. First, the model mismatch between the OVRV and W99 has been appropriately handled by the altuistic controller, showing clear signs of robustness. We can also report some degradation of the performance (compared to the nominal case in the previous subsection).Fig. 5a–b show the benefits of the controller in general, as the mean acceleration is reduced and the mean velocity is increased for all altruism levels in comparison to the coordination control free driving model. Furthermore, the proposed MPC controller indicates lower mean and variance leading to consistency in the driving and reducing oscillations.

Fig. 6a–d depict the cumulative emissions, which suggests that the controller can provide significant reductions in emissions in comparison

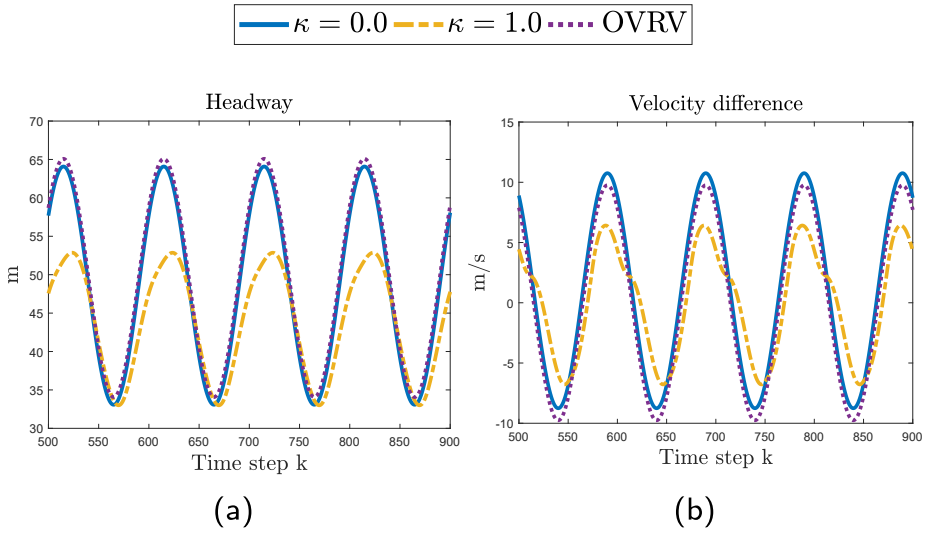

Fig. 3. Headway and velocity difference between the HDV following the CAV for each controller.

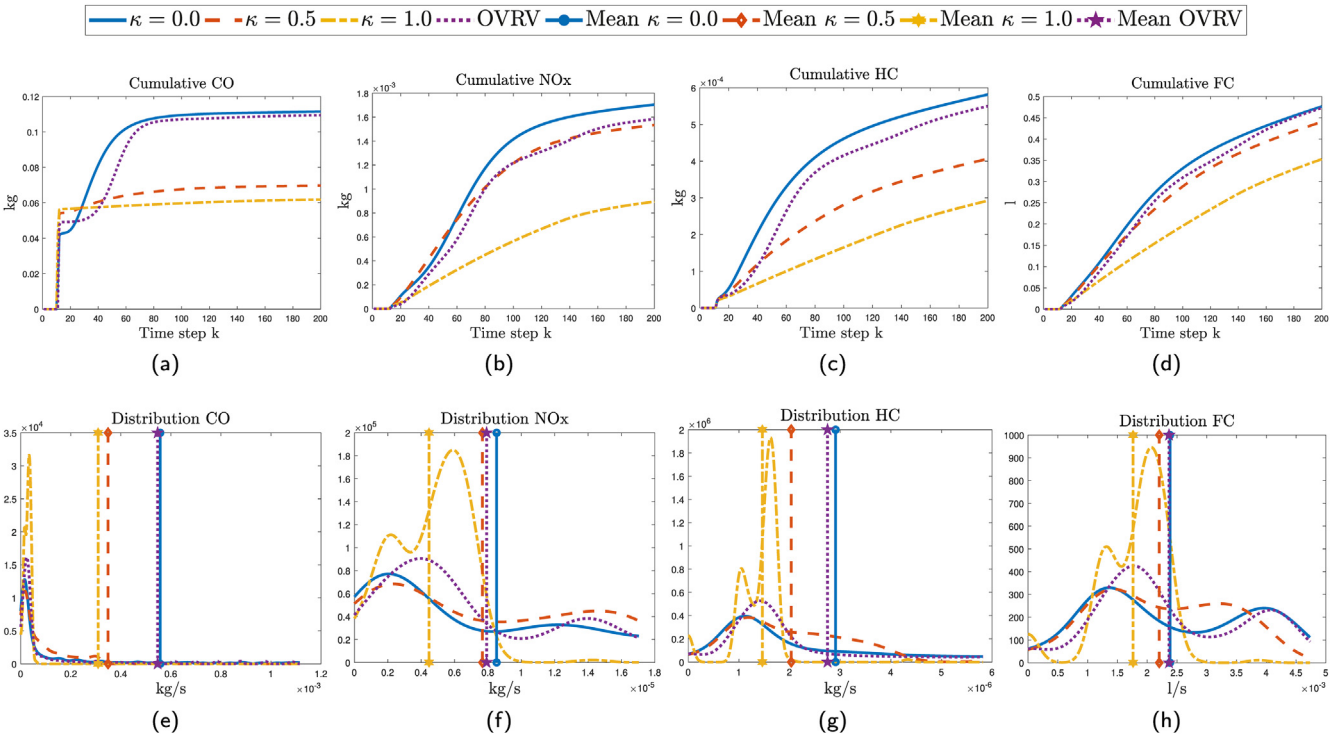

Fig. 4. Single lane OVRV simulation results of vehicle emissions with varying altruism parameter κ.

to W99. However, there is no noticeable difference between the altruism levels. For emission distributions in Fig. 6e–h, the pure altruistic driving strategy exhibits a slight reduction in the variance in comparison to the other altruism levels and the W99 model. Mean values for all altruism are similar and all altruism levels show lower mean values in comparison to uncontrolled W99 models.

5.3.3. *Multi-lane - W99*

For multi-lane simulations, the benefits of altruism are more pronounced compared to single-lane scenarios. These simulations involve the optimization of longitudinal acceleration along with lane changing decisions, meaning that the full control strategy is utilized in this scenario. Furthermore, in this simulation scenario, we are running three CAVs in a decentralized control strategy. In other words, CAVs do not coordinate their actions, and they view the other CAVs as regular HDVs.

In Fig. 7a, we observe that the mean accelerations are reduced by adopting the pure altruistic driving strategy, while the controller in general outperforms the W99 driving model in both mean values and variance. The variance between the different altruism levels is similar. For the velocity distribution in Fig. 7b, all altruism levels and the W99 model perform similarly in terms of mean values; however, the variance differs. Pure altruism exhibits the lowest variance, while the other altruism levels provide slightly higher variance. The W99 model performs the worst when it comes to variance, indicating that the traffic efficiency is negatively impacted.

To investigate vehicle emissions in multi-lane VISSIM scenarios, Fig. 8a–d illustrate the cumulative emissions with respect to time. As seen from Figs. 8a–d, the altruistic driving strategy significantly outperforms

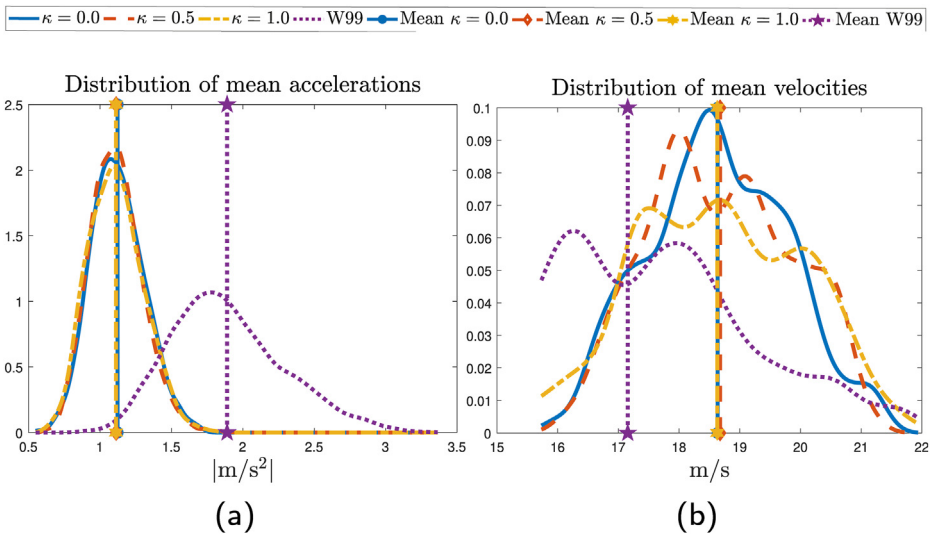

Fig. 5. Single lane VISSIM simulation results of mean acceleration and velocity distribution with varying altruism parameter κ.

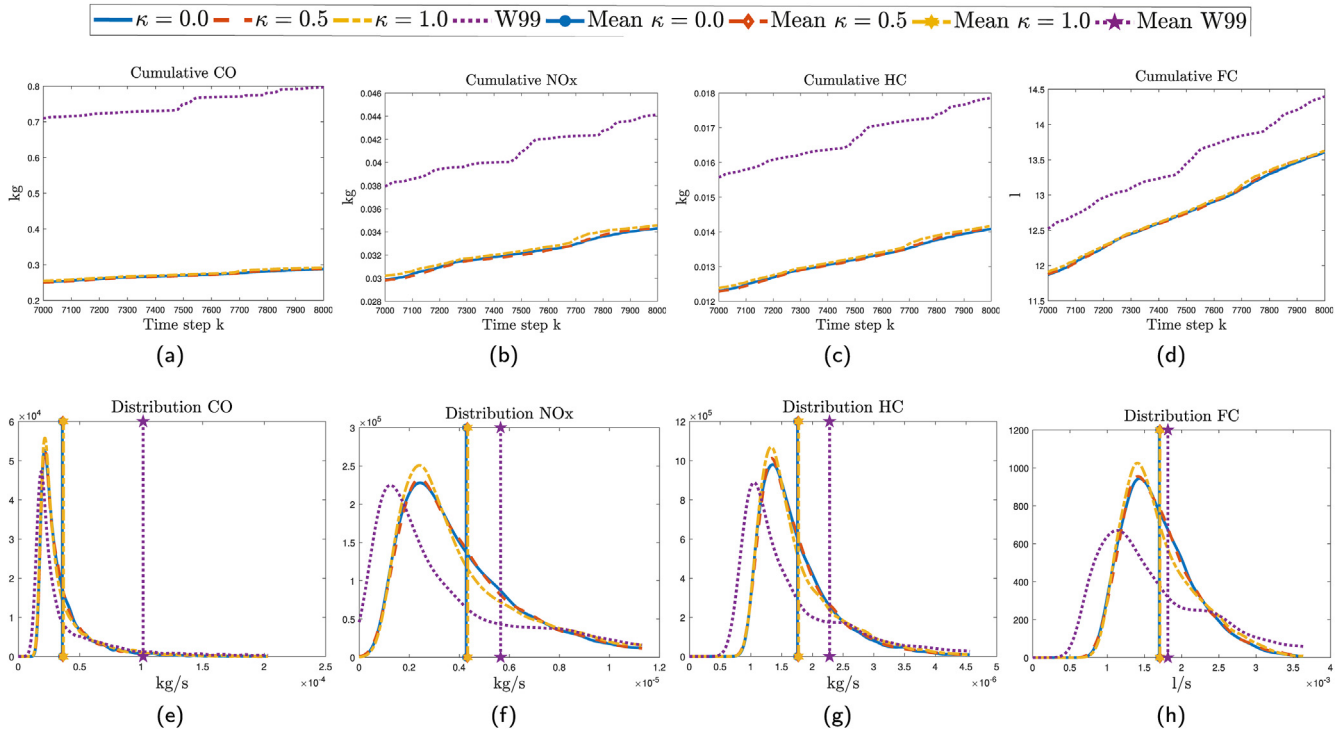

Fig. 6. Single lane VISSIM simulation results of vehicle emissions with varying altruism parameter κ.

the W99 driving model, with the performance gap increasing with higher levels of altruism. In addition, Fig. 8e–h show the histograms of the fuel consumption metrics, which exhibit similar trends, i.e., the pure altruistic driving strategy achieves lower mean emission values than the other controllers.

To explore the CAV-specific results, Fig. 9 shows the distribution of acceleration and velocities for only the CAVs in the simulation. As expected, the highest mean CAV acceleration is obtained in the case of pure altruism (i.e., $\kappa = 1$) since an altruistic CAV sacrifices its own driving objectives for the sake of overall traffic smoothness. The performance difference between the altruistic and selfish controllers in terms of mean acceleration is 20%, while the mean velocities are very close to each other. In addition, the variance in velocity is slightly lower for the altruistic case and the distribution is skewed slightly towards higher velocities as well, giving an indication that the traffic efficiency for HDVs is improved at almost no cost of the efficiency of the CAVs.

Fig. 10 showcases the selfish controller, $\kappa = 0$, altruistic controller $\kappa = 1$ and the penetration rate of CAVs, i.e. the ratio between HDVs and CAVs. For penetration rates between 50% to 25%, the performance is lacking due to the larger mismatch between driving models, which results from decentralized control strategy and that the CAVs assume all vehicles it sees are HDVs. The penetration rate of 20% is the best performing one in terms of mean acceleration, and the penetration rates of 10% and 5% do have lower variance than the rest. This pattern is also present in the velocity distributions, where lower penetration rates show a lower variance meanwhile the mean values for all penetration rates are similar, with a slight increase for the 5% case. This indicates that even a low penetration rate will provide the benefits of altruism to the traffic.

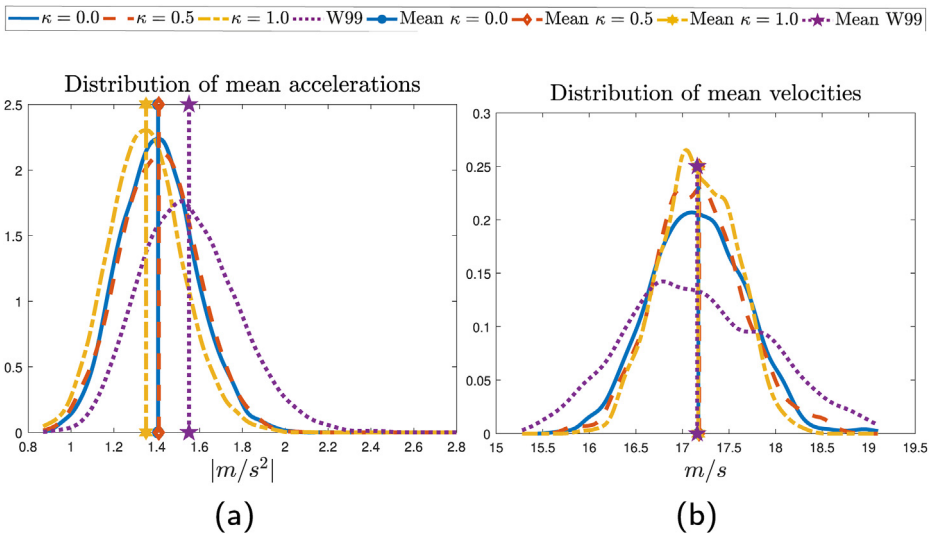

Fig. 7. Multi lane VISSIM simulation results of mean acceleration and velocity distribution with varying altruism parameter κ.

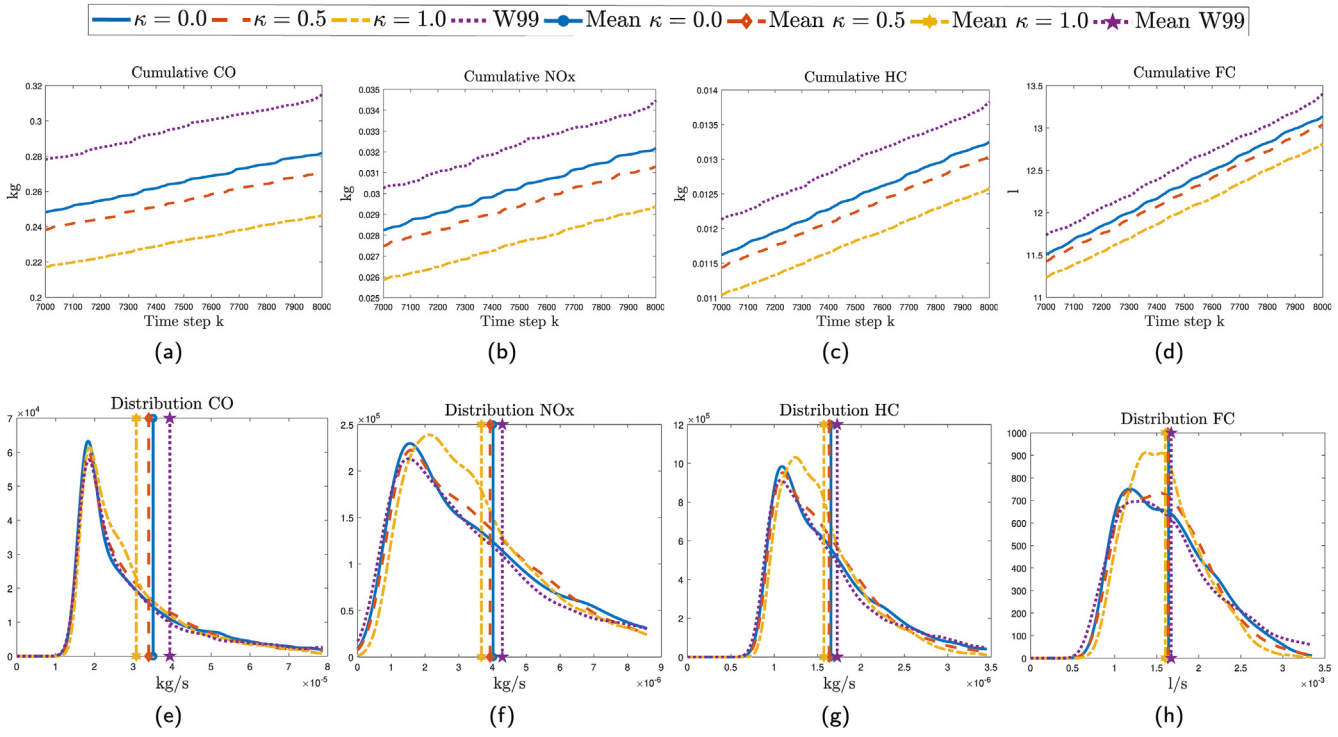

Fig. 8. Multi lane VISSIM simulation results of vehicle emissions with varying altruism parameter κ.

It is worth emphasizing the importance of the model used in predictive control, which influences the performance reached by the altruistic controller. We therefore envisage the benefits of (i) using more sophisticated car-following models in the MPC design or (ii) robustifying the nominal MPC control algorithm, which may help to reach better closed loop performance.

6. Conclusion

We proposed rolling horizon, model based control methodologies to coordinate connected automated vehicles (CAVs) in multi-lane highway traffic conditions. A set of realistic traffic scenarios are defined where automated and human driven vehicles (HDVs) have to co-exist. The objective function of the control algorithms for CAVs is formulated in a such a way that selfish and altruistic goals can be addressed separately or simultaneously. In this way, pro-social human behaviour can be replicated via the control of CAVs. Simulation results in cumulative and empirical distribution metrics showed that traffic efficiency and comfort metrics could be significantly improved by applying an altruistic driving strategy. We pointed out that pro-social behaviour could be triggered with changing the relative weights of the overall cost function, i.e., changing the weightings between the two objectives, namely, the objective function to self and objective function to others.

Future research directions may involve changing the proposed objective functions with the number of vehicles, i.e., traffic density dependent weightings (e.g., no altruism is needed in free-flow conditions).

In addition to the decentralized control strategies proposed in this

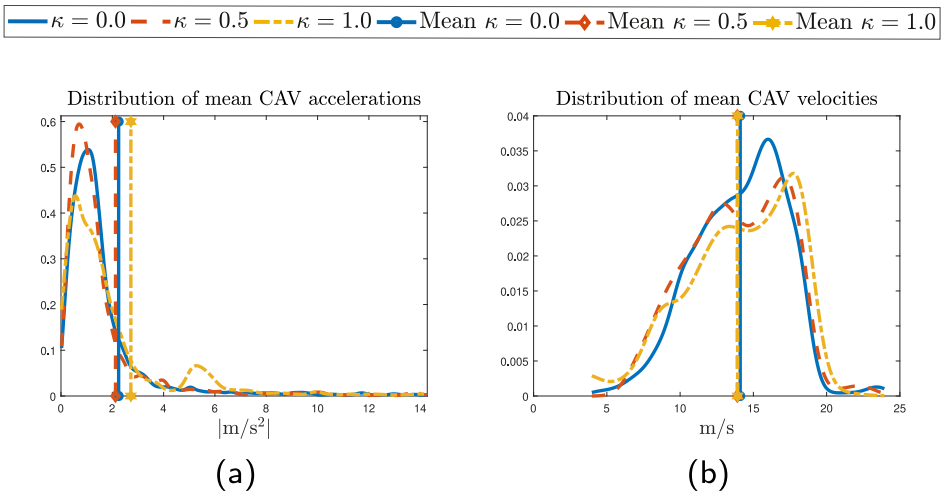

Fig. 9. Multi lane VISSIM simulation results of mean acceleration and velocity distribution with varying altruism parameter κ for the CAVs.

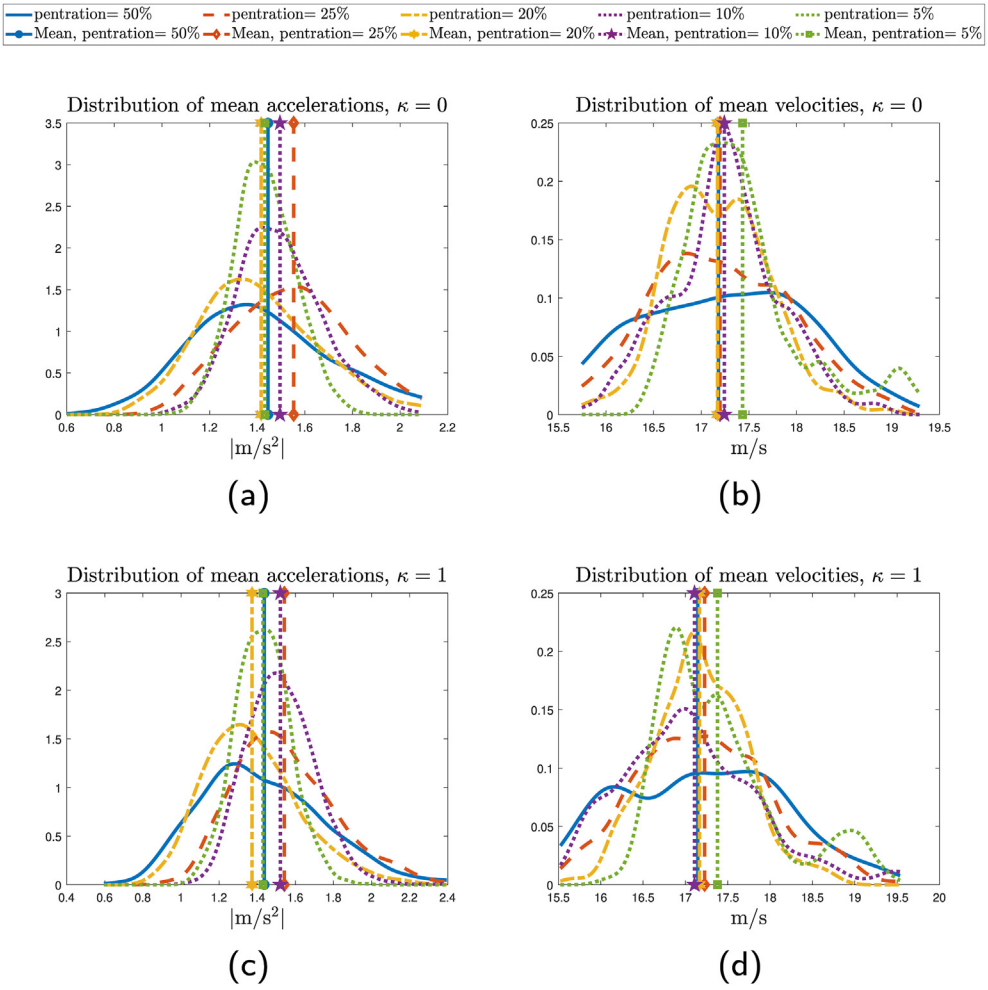

Fig. 10. Multi lane VISSIM simulation results of mean acceleration and velocity distribution with varying CAV penetration rate from 10% to 50%.

paper, more complex MPC controllers utilizing centralized control of CAVs on highways, or cooperative CAV driving, can be designed to incorporate inter-CAV communication for coordinating CAV driving and improved controllability of the traffic flow.

Finally, stability constraints and communication delays are not considered in the proposed predictive control framework. Therefore addressing the stability of adhoc delayed vehicle formations with lane-change maneuvers can be considered as future research direction.

Declaration of competing interest

The authors declare the following financial interests/personal

relationships which may be considered as potential competing interests: Musa Furkan Keskin reports financial support was provided by Transport Area of Advance, Chalmers U university of Technology, Sweden.

Acknowledgements

This work has partially been supported and funded by the Transport Area of Advance. The project IRIS is acknowledged for financial support.

Appendix

A. MPC optimization parameters

The optimization parameters involve the constant terms in the objective function and constraints in Eq. (31) used in the simulation sequences. Since we are evaluating the altruistic parameter, all parameters except κ remain static. The static optimization parameters are set as follows: $\alpha = 2$, $\beta = 2$, $h_{\max} = 70$ m, $h_{\min} = 10$ m, $a_{\max} = 5$ m/s^2, $a_{\min} = -5$ m/s^2, $v_{\max} = 30.5$ m/s, $t_{\min} = 0.25$ s, $w_1 = 0.75$, $w_2 = 0.5$, $\lambda = 0.99$ and $N_p = 40$. For numerical stability in optimization, we perform scaling of the objective function in Eq. (31), as described in the following.

A.1. Scaling the objective function in Eq. (31)

While not critical for the purposes of formulating the MPC controller, there are some practical issues to consider when optimizing the objective function. In order to avoid the practical issues related to differing scales of the numerical components of the objective function, we normalize each objective in Eq. (18), Eq. (19) and Eq. (20) and the slack objective in Eq. (26) to lie within the interval of [0, 1] by dividing each component by its maximum value. Additionally, we also modify the weights w_1, w_2, κ and λ to lie within the same interval. Furthermore, we adopt a decentralized control strategy, i.e., each optimization concerns only a single CAV. However, there may be multiple CAVs on the road, in which case the optimized CAV treats the other CAVs as regular HDVs in predictive optimization.

We modify the efficiency objective in Eq. (18) by

$$\mathcal{J}_i^{\text{eff,scaled}}(\mathbf{u}^{\text{CAV}}_{i,k:k+N_p-1}) = \sum_{n=0}^{N_p-1}\left[\kappa\left(\frac{v^{\text{CAV}}_{k+n,-} - V^*}{v_{\max}}\right)^2 + (1-\kappa)\sum_{j\in\mathcal{G}^{\text{HDV}}_{i,k}}\left(\frac{v^{\text{HDV}}_{j,k+n} - V^*}{v_{\max}}\right)^2\right] \tag{35}$$

where $\{\kappa \in \mathbb{R} \mid 0 \le \kappa \le 1\}$. Similarly, the acceleration magnitude objective Eq. (19) and jerk objective Eq. (20) now becomes

$$\mathcal{J}_i^{\text{mag,scaled}}(\mathbf{u}^{\text{CAV}}_{i,k:k+N_p-1}) = \sum_{n=0}^{N_p-1}\left[\kappa\left(\frac{a^{\text{CAV}}_{i,k+n}}{a_{\max}}\right)^2 + (1-\kappa)\sum_{j\in\mathcal{G}^{\text{HDV}}_{i,k}}\left(\frac{a^{\text{HDV}}_{j,k+n}}{a_{\max}}\right)^2\right] \tag{36}$$

$$\mathcal{J}_i^{\text{jerk,scaled}}(\mathbf{u}^{\text{CAV}}_{i,k:k+N_p-1}) = \sum_{n=0}^{N_p-1}\left[\kappa\left(\frac{a^{\text{CAV}}_{i,k+n+1} - a^{\text{CAV}}_{i,k+n}}{a_{\max}\Delta t}\right)^2 + (1-\kappa)\sum_{j\in\mathcal{G}^{\text{HDV}}_{i,k}}\left(\frac{a^{\text{HDV}}_{j,k+n+1} - a^{\text{HDV}}_{j,k+n}}{a_{\max}\Delta t}\right)^2\right] \tag{37}$$

The objective function concerning the slack variable is also modified as

$$\mathcal{J}^{\text{slack,scaled}} = \sum_{j\in\mathcal{H}^{\text{HDV}}_{i,k}}\left|\frac{\gamma_j}{\max_j \gamma_j}\right|^2 \tag{38}$$

The scaled version of the objective function in Eq. (31) is then given by

$$\begin{aligned}\mathcal{J}_i^{\text{tot,scaled}}(\mathbf{u}^{\text{CAV}}_{i,k:k+N_p-1}) = (1-\lambda)\Big[&(1-w_1)\mathcal{J}_i^{\text{eff,scaled}}(\mathbf{u}^{\text{CAV}}_{i,k:k+N_p-1}) \\ &+w_1\Big((1-w_2)\mathcal{J}_i^{\text{mag,scaled}}(\mathbf{u}^{\text{CAV}}_{i,k:k+N_p-1}) \\ &+w_2\mathcal{J}_i^{\text{jerk,scaled}}(\mathbf{u}^{\text{CAV}}_{i,k:k+N_p-1})\Big)\Big] \\ &+\lambda\mathcal{J}^{\text{slack,scaled}}.\end{aligned} \tag{39}$$

B. Simulation parameters

While the OVRV simulation setup uses the same parameters as the driving model in the MPC formulation, the VISSIM does use a more detailed car-following model (Wiedemann99). The latter model has significantly more parameters: CC0 is the desired vehicle standstill distance, CC1 is the headway time (in seconds) that the vehicle wants to keep. CC0 and CC1 define the safe vehicle headway by $h_{safe} = CC0 + CC1\dot{v}$ where v [m/s] is the vehicle velocity. CC2 controls the variation in meters by defining the oscillation boundaries such as $h_{safe} \geq h \geq h_{safe} + CC2$. CC3 is a threshold parameter; defining the distance when the vehicle in front has been recognized. CC3 is responsible for breaking down and entering into a car-following stage in W99. CC4 and CC5 controls the speed differences during the car-following stage for deceleration and acceleration respectively. CC6 controls the speed dependency of oscillations, where larger values lead to larger velocity with increasing distance while the vehicle is in the car-following stage. CC7, CC8 and CC9 controls the acceleration during oscillation, at standstill, and at above 80 km/h, respectively.

In the VISSIM simulations, the following parameters for W99 Gao (2008) are used, CC0 = 1 [m], CC1 = 0.9 [s], CC2 = 1 [m], CC3 = − 8, CC4 = − 0.05, CC5 = 0.05, CC6 = 1, CC7 = 10 [m/s^2], CC8 = 10 [m/s^2], CC9 = 10 [m/s^2]. Furthermore, the look ahead distance is set to 150 m with 2 vehicles observed at most and the look back distance is set to 100 m. For the single lane scenarios, there are 5 HDVs behind the CAV. For the multi lane scenarios, there are 15 HDVs behind the 3 CAVs.

References

Avedisov, S.S., Bansal, G., Orosz, G., 2020. Impacts of connected automated vehicles on freeway traffic patterns at different penetration levels. IEEE Trans. Intell. Transport. Syst. 1, 1–14.

Bahram, M., 2017. Interactive Maneuver Prediction and Planning for Highly Automated Driving Functions. PhD thesis, PhD Dissertation. Technische Universität München, München, Germany.

Bıyık, E., Lazar, D., Pedarsani, R., Sadigh, D., 2018. Altruistic Autonomy: Beating Congestion on Shared Roads, p. 11978 arXiv1810.

Boyd, S., Vandenberghe, L., 2004. Convex Optimization. Cambridge University Press, Boston, USA.

Buckman, N., Pierson, A., Schwarting, W., Karaman, S., Rus, D., 2019. Sharing is caring: socially-compliant autonomous intersection negotiation. In: Proceedings of 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS, pp. 6136–6143.

Camponogara, E., Zhou, H., Talukdar, S., 2006. Altruistic agents in uncertain dynamic games. J. Comput. Syst. Sci. Int. 45 (4), 536–552.

Csikos, A., Kulcsar, B., 2017. Variable speed limit design based on mode dependent cell transmission model. Transport. Res. C Emerg. Technol. 85, 429–450.

Dabiri, A., Kulcsár, B., 2017. Distributed ramp metering—a constrained discharge flow maximization approach. IEEE Trans. Intell. Transport. Syst. 18 (9), 2525–2538.

Desquesnes, G., Lozenguez, G., Doniec, A., Duviella, E., 2017. Distributed MDP for water resources planning and management in inland waterways. In: Proceedings of 20th IFAC World Congress, 50, pp. 6576–6581, 1.

Dollar, R.A., Vahidi, A., 2018. Efficient and collision-free anticipative cruise control in randomly mixed strings. IEEE Trans. Intell. Veh. 3 (4), 439–452.

Gao, Y., 2008. Calibration and Comparison of the VISSIM and INTEGRATION Microscopic Traffic Simulation Models. PhD thesis, PhD Dissertation. Virginia Tech, USA.

Hamdar, S.H., Treiber, M., Mahmassani, H.S., Kesting, A., 2008. Modeling driver behavior as sequential risk-taking task. Transport. Res. Rec. 2088 (1), 208–217.

Huang, M., Caines, P.E., Malhamé, R.P., 2010. Social dynamics in mean field LQG control: egoistic and altruistic agents. In: Proceedings of 49th IEEE Conference on Decision and Control. CDC), pp. 3140–3145.

Kamal, M.A.S., Imura, J., Hayakawa, T., Ohata, A., Aihara, K., 2014. Smart driving of a vehicle using model predictive control for improving traffic flow. IEEE Trans. Intell. Transport. Syst. 15 (2), 878–888.

Kamal, M.A.S., Taguchi, S., Yoshimura, T., 2016. Efficient driving on multilane roads under a connected vehicle environment. IEEE Trans. Intell. Transport. Syst. 17 (9), 2541–2551.

Keskin, F., Peng, B., Kulcsár, B., Wymeersch, H., 2020. Altruistic control of connected automated vehicles in mixed-autonomy multi-lane highway traffic. In: Proceedings of 2020 IFAC World Congress.

Kesting, A., Treiber, M., Helbing, D., 2010. Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity. Phil. Trans. Math. Phys. Eng. Sci. 368 (1928), 4585–4605.

Kheterpal, N., Parvate, K., Wu, C., Kreidieh, K., Vinitsky, E., Bayen, A., 2018. Flow: deep reinforcement learning for control in sumo. EPiC Ser. Eng. (2), 34–151.

Levy, N., Ben-Elia, E., 2016. Emergence of system optimum: a fair and altruistic agent-based route-choice model. Procedia Comput. Sci. 83, 928–933.

Luspay, T., Kulcsár, B., Varga, I., Zegeye, S., De Schutter, B., Verhaegen, M., 2010. On acceleration of traffic flow. In: Proceedings of 13th International IEEE Conference on Intelligent Transportation Systems, pp. 741–746.

Luspay, T., Péni, T., Kulcsár, B., 2012. Constrained Freeway Traffic Control via Linear Parameter Varying Paradigms. Publisher, Springer, Boston, MA.

M Treiber, A.K., 2013. Modeling human aspects of driving behavior. In: Traffic Flow Dynamics. Springer, Berlin, Germany.

Mavrogiannis, C., DeCastro, J., Srinivasa, S., 2020. Implicit Multiagent Coordinationat Unsignalized Intersections via Multimodalinference Enabled by Topological Braids, 05205 arXiv:2004.

Orosz, G., Wilson, R.E., Stépán, G., 2010. Traffic jams: dynamics and control. Phil. Trans. Math. Phys. Eng. Sci. 364, 4455–4479.

Orosz, G., Wilson, R.E., Szalai, R., Stépán, G., 2009. Exciting traffic jams: nonlinear phenomena behind traffic jam formation on highways. Phys. Rev. E 80, 046205.

Pereira, M., Baykas, P.B., Kulcsár, B., Lang, A., 2022. Parameter and Density Estimation from Real-World Traffic Data: A Kinetic Compartmental Approach. Transportation Research Part B, Methodological.

Pierson, A., Schwarting, W., Karaman, S., Rus, D., 2020. Weighted buffered voronoi cells for distributed semi-cooperative behavior. In: Proceedings of 2020 IEEE International Conference on Robotics and Automation. ICRA, pp. 5611–5617.

Stern, R., Cui, S., Monache, M.D., Bhadani, R., Bunting, M., Churchill, M., Hamilton, N., Haulcy, R., Pohlmann, H., Wu, F., Piccoli, B., Seibold, B., Sprinkle, J., Work, D., 2018. Dissipation of stop-and-go waves via control of autonomous vehicles: field experiments. Transport. Res. C Emerg. Technol. 89, 205–221.

Sugiyama, Y., Fukui, M., Kikuchi, M., Hasebe, K., Nakayama, A., Nishinari, K., Tadaki, S., Yukawa, S., 2008. Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam. New J. Phys. 10 (3), 033001.

Wang, M., Daamen, W., Hoogendoorn, S.P., van Arem, B., 2016. Cooperative car-following control: distributed algorithm and impact on moving jam features. IEEE Trans. Intell. Transport. Syst. 17 (5), 1459–1471.

Wang, N., Wang, X., Palacharla, P., Ikeuchi, T., 2017. Cooperative autonomous driving for traffic congestion avoidance through vehicle-to-vehicle communications. In: Proceedings of 2017 IEEE Vehicular Networking Conference. VNC, pp. 327–330.

Wilson, R.E., Ward, J.A., 2011. Car-following models: fifty years of linear stability analysis–a mathematical perspective. Transport. Plann. Technol. 34 (1), 3–18.

Yu, C., Wang, X., Xu, X., Zhang, M., Ge, H., Ren, J., Sun, L., Chen, B., Tan, G., 2019. Distributed multiagent coordinated learning for autonomous driving in highways based on dynamic coordination graphs. IEEE Trans. Intell. Transport. Syst. 1–14.

Yu, Y., Jiang, R., Qu, X., 2021. A modified full velocity difference model with acceleration and deceleration confinement: calibrations, validations, and scenario analyses. IEEE Intell. Transport. Syst. Mag. 13 (2), 222–235.

Zegeye, S., De Schutter, B., Hellendoorn, J., Breunesse, E., Hegyi, A., 2013. Integrated macroscopic traffic flow, emission, and fuel consumption model for control purposes. Transport. Res. Part C 31, 158–171.

Zhang, L., Orosz, G., 2016. Motif-based design for connected vehicle systems in presence of heterogeneous connectivity structures and time delays. IEEE Trans. Intell. Transport. Syst. 17 (6), 1638–1651.

Zhang, L., Orosz, G., 2018. Beyond-line-of-sight identification by using vehicle-to-vehicle communication. IEEE Trans. Intell. Transport. Syst. 19 (6), 1962–1972.

Jacob Larsson recieved the B.S degree in Automation and Mechatronics, the M.Sc in Product Development and the the M.Sc in Systems, Control and Mechatronics from Chalmers University of Technology, Gothenburg, Sweden, in 2016, 2018 and 2020. He is currently a research engineer at Husqvarna Construction Products Primary R&D, Jonsered, Sweden, within mechatronics and control systems. His main research interests is optimal control systems and modelling.

Musa Furkan Keskin is a researcher and a Marie Skłodowska-Curie Fellow (MSCA-IF) in the department of Electrical Engineering at Chalmers University of Technology, Gothenburg, Sweden. He obtained the B.S., M.S., and Ph.D degrees from the Department of Electrical and Electronics Engineering, Bilkent University, Ankara, Turkey, in 2010, 2012, and 2018, respectively. He received the 2019 IEEE Turkey Best Ph.D Thesis Award for his thesis on visible light positioning systems. His project "OTFS-RADCOM: A New Waveform for Joint Radar and Communications Beyond 5G″ is granted by the European Commission through the H2020-MSCA-IF-2019 call. His current research interests include intelligent transportation systems, joint radar-communications, and positioning in 5G and beyond 5G systems.

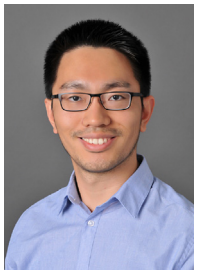

Bile Peng received the B.S. degree from Tongji University, Shanghai, China, in 2009, the M.S. degree from the Technische Universität Braunschweig, Germany, in 2012, and the Ph.D. degree with distinction from the Institut für Nachrichtentechnik, Technische Universität Braunschweig in 2018. He has been a Postdoctoral researcher in the Chalmers University of Technology, Sweden from 2018 to 2019, a development engineer at IAV GmbH, Germany from 2019 to 2020. Currently, he is a Postdoctoral researcher in the Technische Universität Braunschweig, Germany. His research interests include wireless channel modeling and estimation, Bayesian inference as well as machine learning algorithms, in particular deep reinforcement learning, for resource allocation of wireless communication. Dr. Peng is a major contributor to the IEEE Standard for High Data Rate Wireless Multi-Media Networks Amendment 2: 100 Gb/s Wireless Switched Point-to-Point Physical Layer (IEEE Std 802.15.3d-2017) and received the IEEE vehicular technology society 2019 Neal Shepherd memorial best propagation paper award.

Balázs Kulcsár received the M.Sc. degree in traffic engineering and the Ph.D. degree from Budapest University of Technology and Economics (BUTE), Budapest, Hungary, in 1999 and 2006, respectively. He has been a Researcher/Post-Doctor with the Department of Control for Transportation and Vehicle Systems, BUTE, the Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN, USA, and with the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands. He is currently a Professor with the Department of Electrical Engineering, Chalmers University of Technology, Göteborg, Sweden. His main research interest focuses on traffic flow modeling and control.

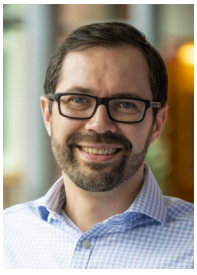

Henk Wymeersch (S′01, M′05, SM′19) obtained the Ph.D. degree in Electrical Engineering/Applied Sciences in 2005 from Ghent University, Belgium. He is currently a Professor of Communication Systems with the Department of Electrical Engineering at Chalmers University of Technology, Sweden. He is also a Distinguished Research Associate with Eindhoven University of Technology. Prior to joining Chalmers, he was a postdoctoral researcher from 2005 until 2009 with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. Prof. Wymeersch served as Associate Editor for IEEE Communication Letters (2009–2013), IEEE Transactions on Wireless Communications (since 2013), and IEEE Transactions on Communications (2016–2018). During 2019–2021, he is a IEEE Distinguished Lecturer with the Vehicular Technology Society. His current research interests include the convergence of communication and sensing, in a 5G and Beyond 5G context.