Provably Safe and Smooth Lane Changes in Mixed Traffic

Maximilian Naumann[1,2], Hendrik Königshof[1,2] and Christoph Stiller[1,2]

***Abstract*— While lane change behavior of human drivers has already been widely investigated, concepts and algorithms to plan lane changes for automated vehicles become necessary through the progress of automated driving. In this paper, we present an approach for safe and smooth lane changes in dense traffic, that consider both the cooperation of humans but also the increased need for safety compared to lane changes performed by human drivers. This real time capable approach is based on path-velocity decomposition, the intelligent driver model, a model compliance check for human machine cooperation and the concept of responsibility sensitive safety. The work is targeted towards mandatory lane changes determined by an upstream routing module, but is also applicable to optional lane changes.**

***Index Terms*— Motion Planning, Trajectory Planning, Behavior Generation, Decision Making, Lane Change, Mixed Traffic.**

I. Introduction

The analysis of human lane change behavior in traffic already started decades ago for the sake of traffic flow simulation. In recent years, advanced driver assistance systems (ADAS) as well as fully automated vehicles are becoming more and more successful [1]. In the development of automated vehicles, motion planning, which is often decomposed into decision making and trajectory planning, is a crucial part.

Motion planning algorithms for single lane driving can mostly be covered in a reactive way. The combination of currently commercially available adaptive cruise control (ACC) and lane keeping assistant (LKA) often already suffice in single lane driving, while safety must still be ensured by the driver, currently. Performing lane changes, however, requires the assumption of cooperation between drivers, at least in dense traffic. Also, lane changes are safety critical maneuvers even for humans, i.a. through the blind spot. Moreover, while lane changes in highway traffic can be considered rather optional and can be performed only in huge gaps, they are often mandatory in urban traffic due to the course of the route. Thus, motion planning of lane changes is a challenging but necessary task in the development of fully automated vehicles. It requires assumptions considering the cooperation of other road traffic participants, but at the same time must be provably safe, in order to never cause a collision. Thus, motion planning and prediction cannot be decoupled anymore, but must be considered in an integrated approach.

[1]FZI Research Center for Information Technology, Mobile Perception Systems, 76131 Karlsruhe, Germany. {naumann, koenigshof}@fzi.de

[2]Karlsruhe Institute of Technology (KIT), Institute of Measurement and Control, 76131 Karlsruhe, Germany. stiller@kit.edu

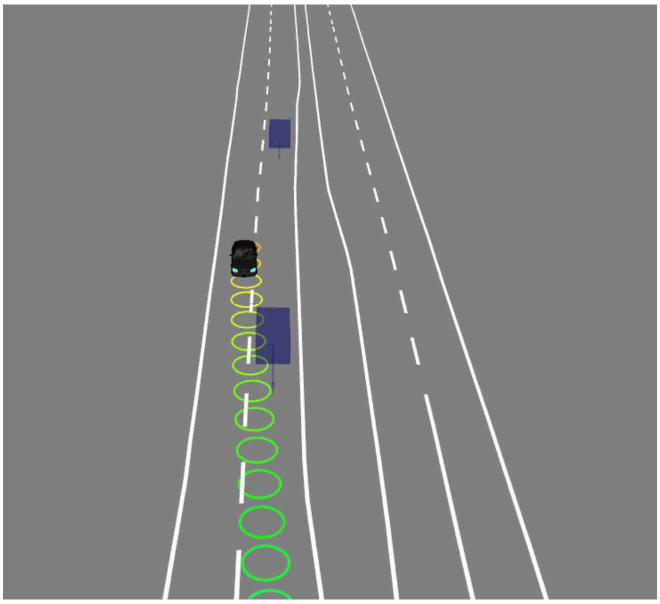

Fig. 1: An exemplary safe lane change trajectory. Key to facilitating safety is leveraging the fast reaction time of automated vehicles, which is why the distance to the leading vehicle is significantly smaller than the distance to the following vehicle.

In this work, we shortly review existing approaches to lane change motion planning for automated vehicles, before we propose a new approach for safe and smooth lane changes in mixed traffic. Mandatory lane changes are the focus, while the approach is also applicable to optional lane changes, if the lane change desire and its possible scope is determined upstream. When introducing the problem statement, we pay special attention to the safety requirements, based on the work of Shalev-Shwartz et al. [2]. In case a lane change requires cooperation of other traffic participants, their compliance is checked with a cooperation model according to our previous work [3].

The contribution of this paper is twofold:

- We propose a formalization of safe and "non-reckless" lane changes as addition to and based on the concept of Responsibility-Sensitive Safety (RSS) [2] and Set-Based Safety Verification [4], [5]
- We propose a general approach to perform these safe lane changes in a comfortable way

The remainder of this paper is structured as follows: Related work is presented in section II. Subsequently, our approach is presented in section III, before we evaluate it in section IV.

II. Related Work

Various approaches exist to determine whether or not an optional lane change is currently beneficial. Kesting et al. define the lane change model MOBIL based on the Intelligent Driver Model (IDM) [6]: They determine the possible current acceleration according to the IDM and decide in favor of the maximum acceleration, while using a switching threshold to prevent lane changes for a marginal advantage [7]. Mandatory lane changes are forced by a virtual vehicle at the end of the current lane. With the latter, however, lane changes are not performed in a foresighted manner, but only when the need to stop for the virtual vehicle is imminent. The particular trajectory for lane change itself is not further considered. Eggert et al. present a similar approach based on their Foresighted Driver Model (FDM) [8], also focusing on the lane change decision [9]. Compared to MOBIL, they can perform foresighted lane changes and lane changes due to tailgating drives, as they consider a potential risk along the possible future trajectories. However, as they predict other traffic participants before considering their own lane change desire, merging into narrow gaps is not possible with their approach.

While the previous approaches focused on driver modeling, other approaches target the development of algorithms to facilitate lane changes in automated vehicles. If the need for a lane change was determined, the lane change process is often divided into gap selection, gap approach, gap evaluation and the subsequent merge into the gap [10].

Pek et al. present an approach to analyze whether a gap is large enough to safely perform a lane change by using set-based reachability analysis [5]. The approach is, to the best of the authors knowledge, the first to prove the safety of lane change maneuvers. They assume that following vehicles in the target lane only have to maintain a safe distance to the changing vehicle once the latter is fully within the target lane. As they further assume that following vehicles potentially fully accelerate up to some threshold over the maximum allowed velocity, lane changes in dense traffic are not possible with their approach. In their application of this verification, they assume a fixed prediction of the respective vehicle, such that their safe distance only covers emergency braking of all participants to any object that is at least partially within their lane [11]. Uncompliant behavior such as trying to avoid the lane change by tailgating or delayed reaction to lane changers is not considered. Therefore, this approach is not suitable for mixed traffic with potentially uncompliant human drivers. Mirchevska et al. are using the same safety verification method for a reinforcement learning based lane change decision making, however they target towards optional lane changes in highway traffic and not towards lane changes in dense traffic [12].

Zhan et al. propose a more general motion planning framework which also includes lane changes [13]. Similar to the work of Miller et al. [11], they assume a known prediction for other obstacles and their safety consideration focuses on potential emergency braking to any object that is at least partially within the lane, but not on uncompliant traffic participants or delayed reaction to lane changers. The previously presented trajectory planning algorithm of Nilsson et al. [14] also facilitates lane changes, but also based on a fixed prediction, neglecting interaction.

This work focuses on mandatory lane changes, i.e. the need for a lane change and the scope in which it can be performed is determined upstream in a routing module. This goal is also targeted by Cunningham et al. [15] and Hubmann et al. [10]. For this purpose, the interaction between vehicles must be considered in the approach, as it is crucial to facilitate lane changes in dense traffic. Assuming fixed prediction often would imply, that lane changes are not feasible at all. For this interaction, more specifically for other traffic participants opening a gap, the ego vehicle relies on others' courtesy. As it is not known whether or not certain traffic participants are going to show courteous behavior or not, Cunningham et al. [15] and Hubmann et al. [10] formulate the problem as a Partially Observable Markov Decision Process. Cunningham et al. simplify this POMDP by using Multi Policy Decision Making (MPDM) and using deterministic transitions

$$P_a(s,s') = \begin{cases} 1, & \text{for } s' = f(s,a) \\ 0, & \text{else} \end{cases}, \tag{1}$$

i.e. simulating a discrete set of policies and choosing the best policy. For the lane change scenario, one of the policies is to attempt to merge into a gap which is too small, and hoping that the follower in the target lane will behave courteous. As in their simulation, all policies avoid collisions by slowing down, their attempts are always successful if feasible. On the other hand, this could also lead to reckless cut-ins just in front of others, which is not suitable for mixed traffic.

Hubmann et al. focus on the POMDP formulation [10]. In addition to the observable position (in lane based Frenet-Serret frames relative to the relevant intersection) and velocity of all vehicles, other (i.e. non-ego) vehicles have a binary hidden state, describing whether they react cooperatively or not. When solving the POMDP, in the forward simulation, the actions of other vehicles are determined by the IDM, either to the front vehicle, or to the merging vehicle, depending on a newly introduced *yield classification* and the previously mentioned hidden state. A vehicle reacts cooperatively if its hidden state is set accordingly and the continuous yield classifier outputs a value above a certain threshold. Provable safety is not focused on in their work. Given a good yield classifier, the agent could even exploit the desire of other vehicles to stay collision free, and determine reckless merges as beneficial.

III. Proposed Approach

The goal of this work is to facilitate lane changes for automated vehicles upon request of their upstream routing module. As safety is the most important criterion for every motion planner, it is a central part in our approach, as explained in the following.

A. Safety Consideration

The safety considerations undertaken in this work are based on the concept of Responsibility Sensitive Safety (RSS) [2] and the concept of Pek et al. [5]. In consensus with both works, we make use of the road structure and define safety in a lane-based way. Neglecting lane changes of other cars, the safety of a lane change depends on the safety with respect to four vehicles: the leading and the following vehicle in both the current (source) and the target lane. In consensus with Pek, the responsibility to maintain a safe distance to following vehicles is considered to be entirely with the ego vehicle. Obviously, the ego vehicle must ensure safety distances to the leader in at least one of the lanes at every point in time. While approaching the gap and preparing the lane change in the source lane, the safe distance to the leader in the source lane must be maintained. While being entirely within the source lane, the follower in the source lane is expected to maintain a safe distance to us.

Further, according to RSS "common sense" rule 5, the ego vehicle shall not cut-in recklessly. Thus, a lane change cannot be performed with an arbitrarily small distance to the follower in the target lane. Pek et al. [5] propose that the follower in the target lane must maintain a safe distance towards the lane changing vehicle as soon as the latter has completed the lane change. Assuming a lane change duration of 4 seconds, and that the follower in the target lane fully accelerates up to some threshold over the maximum allowed velocity during this time, huge gaps would be needed for safe lane changes. Instead, we propose the following:

A cut-in is not considered reckless with respect to a following vehicle in the target lane, if

- it has been indicated using the "direction-indicator lamps"
- it has been started with at least the safe distance w.r.t. to the following vehicle
- after slowly moving across the border of the target lane, the following vehicle in the target lane has maintained a safe distance during a lane change response time δ_{LCR}.

Still, the lane change might need to be abandoned due to tailgating of the following vehicle in the target lane, or due to an emergency brake of either of the leading vehicles. In order to maintain safety along the source lane, the following vehicle must keep maintaining a safe distance to us. For this purpose, we propose:

A following vehicle has to maintain a safe distance to vehicles that are ahead and leaving the lane until they have fully left the lane. Additionally, human drivers might assume, that the lane change will no longer be abandoned at some point during its execution. Thus, it might be allowed to go below this longitudinal safe distance if the lane changing vehicle has left the lane with for example $2/3$ of their geometry and a lateral safe distance is maintained. From the ego vehicle's point of view, this causes a more conservative behavior: In order to keep safe in the source lane, the ego vehicle must stay mainly within this lane during the lane change response time δ_{LCR}.

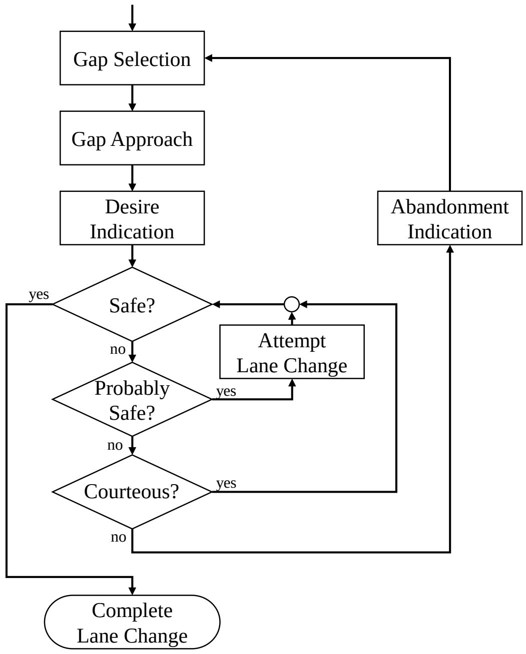

Fig. 2: Lane change process overview. The detailed procedure is explained in the text.

Consequently, starting from a safe situation in the source lane, a safe lane change can be performed by

- approaching a state in the source lane that would be considered longitudinally safe w.r.t. all vehicles in the target lane, while maintaining safety in the source lane
- indicating the lane change
- slowly moving across the border of the target lane, while maintaining a safe distance to the follower and leader in the target lane, and to the leader in the source lane (*)
- maintaining these safe distances for at least a lane change response time δ_{LCR} (*)
- completing the lane change while maintaining a safe distance to the leader in the target lane

while the lane change can be safely abandoned in the stages marked with an asterisk (*). Note that, as motivated by the RSS concept [2], a *safe* lane change is guaranteed not to *cause* a collision according to the previously defined rules, while a single vehicle cannot ensure that it will never be involved in a collision. The reader is referred to [2] for details.

B. Lane Change Process Overview

Motivated by the safety consideration, the lane change process is designed as follows: Once the lane change desire has been determined by the routing module, possible gaps along the target lane are analyzed and the most promising gap is chosen. Subsequently the gap is approached and the lane change desire is communicated to the other traffic participants. Then, a safety check is performed, in order to determine whether or not a lane change is safe already, will probably be safe in the future or whether courtesy of the

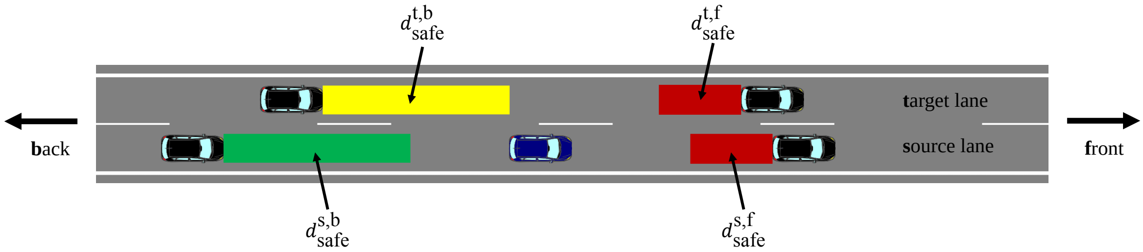

Fig. 3: Safe distances in the lane change process. While the distance visualized in green is to be maintained by the other traffic participant, the ego vehicle is responsible to maintain the red distances throughout the lane change. The yellow distance is to be maintained by the ego vehicle only for a certain lane change response time δ_{LCR}, as further explained in the text.

following vehicle in the target lane is needed. In case the lane change is safe, it is completed. In case courtesy is needed, the behavior of the potentially courteous vehicle is continuously monitored (compliance check) in order to decide whether to abandon the lane change into this gap or to hope for courtesy, i.e. the gap becoming larger. The overview is visualized in Fig. 2.

C. Lane Change Preparation

The gap analysis and selection is not the focus of this work. It can be performed using either a hand tuned heuristic, or machine learning, on either an abstract state space, or even raw sensor data. For mandatory lane changes, it is certainly advisable to attempt the lane change as early as possible, to have more options remaining in case an attempt fails. This idea can easily be implemented in a heuristic, but will probably also be the outcome of a learned value function. A simple heuristic that was used for the evaluation of our approach is presented in section IV.

Having chosen a gap, the goal is to perform the approach fast and comfortably while maintaining the safe distances in the source lane and trying to reach safe distances also in the target lane. Note that this approach can be performed by catching up to a gap ahead or drop back to a gap behind, if permitted by the traffic rules. Having approached the gap, indicating the lane change desire via the direction-indicator lamps is obvious.

D. Lane Change Execution

After indicating the lane change desire, the current distances to the following and the leading vehicle in the target lane have to be investigated. As the leading vehicle in the target lane is not expected to react to the ego vehicle, a safe distance to it should be kept throughout the *safety approval*, as long as the ego vehicle is still ahead of the following vehicle when keeping this distance. For the latter distances, the following conditions are distinguished, using the safe distances from Fig. 3:

C1 Complete lane change is possible: The gap size is large enough, such that a safe distance is maintained throughout the complete lane change, which takes a time $\Delta t_{\mathrm{LC}} \geq \delta_{\mathrm{LCR}}$, even if the following vehicle fully accelerates up to $v_{\max}$ and the lead vehicle fully brakes. For the calculation of the necessary distance $d^{\mathrm{t,b}}_{\mathrm{safe,clc}}$, which is significantly larger than $d^{\mathrm{t,b}}_{\mathrm{safe}}$, it is referred to the work of Pek et al. [5].

C2 Safe distance is currently fulfilled: The gap size is large enough, such that currently both the safe distance to the leader and to the follower is maintained. For the calculation of the distances $d^{\mathrm{t,f}}_{\mathrm{safe}}$ and $d^{\mathrm{t,b}}_{\mathrm{safe}}$, it is referred to work about RSS [2].

C3 Safe distance is probably fulfilled at time t^*: The gap size is large enough, such that both the safe distance to the leader and to the follower will be maintained at time t^* (assuming the current motion plan of the ego vehicle and the current prediction of others). For the calculation of the distances $d^{\mathrm{t,f}}_{\mathrm{safe}}(t^*)$ and $d^{\mathrm{t,b}}_{\mathrm{safe}}(t^*)$, it is referred to the concept of RSS [2], interpolating the states at t^* using the prediction and the current ego vehicle motion plan.

With these cases, the behavior is defined as follows: If condition C1 is fulfilled, the lane change is performed immediately. However, this condition is most likely not fulfilled in normal road traffic, as it implies huge gaps.

If C1 is not fulfilled, the ego vehicle needs to make use of the possibility to approve its lane change via moving across the border of the target lane and granting a response time δ_{LCR}. Assuming it could reach the border of the target lane at time $t_b = t_0 + \Delta t$, i.e. some time Δt after the current time t_0, and using a non-cooperative[1] prediction. If conditions C2 is fulfilled and C3 is fulfilled for $t^* \in [t_0, t_b + \delta_{\mathrm{LCR}}]$, this means that the lane change approach will probably be successful. In this case, the approach is pursued and the ego vehicle moves towards the target lane. If the safe distance is maintained by the following vehicle, the lane change will be completed after δ_{LCR} has passed. Otherwise, the ego vehicle returns to the source lane[2].

If conditions C2 and C3 are not fulfilled given the previous assumptions, the ego vehicle requires cooperation of the follower, which is called courtesy in this case, as the follower does at least not benefit immediately. Thus, a cooperative prediction model must be used, and it must be checked whether the vehicle behaves compliant to the cooperation model. In case it does, conditions C2 and C3 will eventually become fulfilled and the lane change can be attempted. In

[1]But also not malicious.

[2]This means that the following vehicle behaved malicious or the prediction model was mistaken.

case it does not, the lane change has to be started over again, selecting a new gap.

IV. Evaluation

In order to show the potential of our approach, a basic implementation in C++ was realized and evaluated it in the simulation framework CoInCar-Sim [16].

A. Implementation

The implementation is based on the Path-Velocity Decomposition as introduced by Kant and Zucker [17]. The reference paths are computed as the centerline of the lanes, using the *lanelet2* map format and library [18]. As trajectory controller, the decoupled approach presented by Ziegler et al. [19] was used.

The gap approach is implemented as follows: The desired position within the gap is determined as $\min(2 \cdot d_{\text{safe}}, \Delta s_{\text{gap}})$, i.e. if possible well behind the safe distance, but at least in front of the following car. The following and the leading vehicle are predicted with constant velocity, thus the gap also moves with constant velocity. This desired position is now approached in a sequence of constant acceleration, constant velocity and again constant acceleration. This is, if the gap is currently ahead[3], it is accelerated up to a certain velocity v^{+}_{peak} to catch up, before decelerating into the gap. Otherwise, it is decelerated up to a certain velocity v^{-}_{peak} to drop back before accelerating into the gap. The gap selection heuristic is to simply take the gap that can be reached earliest in space along the target lane.

Having approached the gap, a longitudinal trajectory is planned within the upper distance of $2 \cdot d_{\text{safe}}$ while assuming constant velocity for the leader. Additionally, the safety distances checks are started. If condition C1 is fulfilled, a lane change path is planned from the source to the target lane using the sigmoid function $\tanh$. Having reached the gap in the target lane, the lane change task is fulfilled and a safe transition to any other planner is feasible.

In the likely case that condition C1 is not fulfilled, an approaching path just across the border of the target lane is planned, using the $\tanh$ function. Conditions C2 and C3 are checked for a non-cooperative constant velocity prediction of the following vehicle. If they are fulfilled, the lane change is completed after having been slightly across the target lane for δ_{LCR} by planning the lane change path as previously described and pursuing it immediately.

If courtesy is required, as explained in section III-D, it is assumed that the follower drives as described by the enhanced IDM and regards the ego vehicle as its front vehicle. The model compliance is checked by a simple time threshold within which the gap is assumed to be opened large enough.

B. Exemplary Scenario

We show the performance of the previously described implementation in a mandatory lane change at Haid-und-Neu-Straße in Karlsruhe, Germany. As always with cooperative

[3]Regarding position and velocity difference.

TABLE I: IDM Parameters.

Desired speed v_{des}	$50\frac{\text{km}}{\text{h}}$	Max. acceleration a	$1.4\frac{\text{m}}{\text{s}^2}$
Desired time gap T	0.7s	Desired deceleration b	$2\frac{\text{m}}{\text{s}^2}$
Jam distance s_0	2m	Coolness factor c	0.99
Acceleration exponent δ	4		

TABLE II: RSS Parameters.

Ego response time δ_{ego}	0.1s
Others' response time δ_{other}	1s
Lane change response time δ_{LCR}	2s
Max. acceleration during response time $a_{\text{max,accel,ego}}$	$2\frac{\text{m}}{\text{s}^2}$
Max. acceleration during response time $a_{\text{max,accel,other}}$	$3\frac{\text{m}}{\text{s}^2}$
Max. deceleration of front vehicle $a^{\text{f}}_{\text{max,brake}}$	$8\frac{\text{m}}{\text{s}^2}$
Min. deceleration of back vehicle $a^{\text{b}}_{\text{min,brake}}$	$7\frac{\text{m}}{\text{s}^2}$

approaches, their success strongly depends on whether the assumptions regarding the cooperation of other vehicles are met or not. For this reason, we refrain from plotting velocity or acceleration profiles, but instead advise a qualitative assessment of the attached video: In case the relevant driver is correctly assumed to be cooperative, the ego vehicle can safely merge. Otherwise, after some time, the ego vehicle decides to abort the merge maneuver into this gap.

Being cooperative or non-cooperative was realized using the enhanced IDM [20] with the parameters given in table I. In the cooperative case, the following vehicle considers the ego vehicle as its front vehicle, in the non-cooperative, it keeps regarding the leading vehicle as its front vehicle. However, we want to emphasize that the previously mentioned success of the approach, that is commonly evaluated in exemplary scenarios, only regards the comfort and the convenience of the executed motion in certain scenarios, while their safety is guaranteed w.r.t. the defined assumptions, following the concept of RSS [2]. The safe distances are computed using the parameters given in table II. Note that due to the short reaction time of the ego vehicle, the safe distances to the front are significantly smaller, which is why the ego vehicle approaches the leading vehicle closely. In addition to the obvious intention communication through moving towards the target lane, being closer to the leading vehicle facilitates smooth lane changes, as the following vehicle does not need to decelerate much.

V. Conclusions and Future Work

In this paper, a novel approach for safe and smooth lane changes in mixed traffic is presented. For this, a formalization of "non-reckless" lane changes by allowing an additional reaction time for vehicles that are affected by a cut-in is introduced. Subsequently, the lane change approach is proposed: It is based on first reaching the gap in the source lane and possible slightly moving towards the target lane, before starting the actual merge into the target lane, which is only performed if the maneuver is provably safe. This approach can be realized in a computationally cost-efficient way, so that it could even be implemented in automobiles with ADAS on board using current series hardware. An exemplary

implementation is shown in the Video Attachment. While the approach is suitable for large gaps, it is also targeted towards dense traffic where courtesy of other traffic participants is needed to facilitate the lane change.
Future work includes testing the approach in our test vehicle *BerthaOne* and analyzing large trajectory datasets to suggest appropriate values for safety parameters.

Acknowledgements

The authors would like to thank Johannes Beck for his valuable suggestions and the fruitful discussions. Furthermore, we gratefully acknowledge support of this work by the Tech Center a-drive, the Profilregion Mobilitätssysteme Karlsruhe and by the Deutsche Forschungsgemeinschaft (German Research Foundation) within the Priority Program "SPP 1835 Cooperative Interacting Automobiles".

References

[1] K. Bengler, K. Dietmayer, B. Färber, M. Maurer, C. Stiller, and H. Winner, "Three decades of driver assistance systems - review and future perspectives", *IEEE Intell. Transp. Syst. Mag.*, vol. 6, no. 4, pp. 6–22, 2014.

[2] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "On a formal model of safe and scalable self-driving cars", *ArXiv preprint arXiv:1708.06374*, 2017.

[3] M. Naumann, M. Lauer, and C. Stiller, "Generating comfortable, safe and comprehensible trajectories for automated vehicles in mixed traffic", in *Int. Conf. on Intell. Trans. Syst. (ITSC)*, Nov. 2018, pp. 575–582.

[4] M. Althoff and J. M. Dolan, "Online verification of automated road vehicles using reachability analysis", *IEEE Transactions on Robotics*, vol. 30, no. 4, pp. 903–918, Aug. 2014.

[5] C. Pek, P. Zahn, and M. Althoff, "Verifying the safety of lane change maneuvers of self-driving vehicles based on formalized traffic rules", in *2017 IEEE Intell. Vehicles Symposium (IV)*, Jun. 2017, pp. 1477–1483.

[6] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations", *Phys. Rev. E*, vol. 62, pp. 1805–1824, 2 Aug. 2000.

[7] A. Kesting, M. Treiber, and D. Helbing, "General lane-changing model mobil for car-following models", *Transportation Research Record*, vol. 1999, pp. 86–94, Jan. 2007.

[8] J. Eggert, F. Damerow, and S. Klingelschmitt, "The foresighted driver model", in *IEEE Intell. Vehicles Symposium (IV)*, Jun. 2015, pp. 322–329.

[9] J. Eggert and F. Damerow, "Complex lane change behavior in the foresighted driver model", in *IEEE Int. Conf. on Intell. Trans. Syst.*, Sep. 2015, pp. 1747–1754.

[10] C. Hubmann, J. Schulz, G. Xu, D. Althoff, and C. Stiller, "A belief state planner for interactive merge maneuvers in congested traffic", in *Int. Conf. on Intell. Trans. Syst. (ITSC)*, Nov. 2018, pp. 1617–1624.

[11] C. Miller, C. Pek, and M. Althoff, "Efficient mixed-integer programming for longitudinal and lateral motion planning of autonomous vehicles", in *2018 IEEE Intell. Vehicles Symposium (IV)*, Jun. 2018, pp. 1954–1961.

[12] B. Mirchevska, C. Pek, M. Werling, M. Althoff, and J. Boedecker, "High-level decision making for safe and reasonable autonomous lane changing using reinforcement learning", in *2018 21st International Conf. on Intell. Trans. Syst. (ITSC)*, Nov. 2018, pp. 2156–2162.

[13] W. Zhan, J. Chen, C. Chan, C. Liu, and M. Tomizuka, "Spatially-partitioned environmental representation and planning architecture for on-road autonomous driving", in *IEEE Intell. Vehicles Symposium (IV)*, Jun. 2017, pp. 632–639.

[14] J. Nilsson, M. Brännström, J. Fredriksson, and E. Coelingh, "Longitudinal and lateral control for automated yielding maneuvers", *IEEE Transactions on Intell. Trans. Syst.*, vol. 17, no. 5, pp. 1404–1414, May 2016, ISSN: 1524-9050.

[15] A. G. Cunningham, E. Galceran, R. M. Eustice, and E. Olson, "MPDM: Multipolicy decision-making in dynamic, uncertain environments for autonomous driving", in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, May 2015, pp. 1670–1677.

[16] M. Naumann, F. Poggenhans, M. Lauer, and C. Stiller, "Coincar-sim: an open-source simulation framework for cooperatively interacting automobiles", in *IEEE Intl. Conf. Intell. Vehicles*, Jun. 2018.

[17] K. Kant and S. W. Zucker, "Toward efficient trajectory planning: the path-velocity decomposition", *Int. J. Rob. Res.*, vol. 5, no. 3, pp. 72–89, 1986.

[18] F. Poggenhans, J.-H. Pauls, J. Janosovits, S. Orf, M. Naumann, F. Kuhnt, *et al.*, "Lanelet2: A high-definition map framework for the future of automated driving", in *Proc. IEEE Intell. Trans. Syst. Conf.*, Hawaii, USA, Nov. 2018.

[19] J. Ziegler, P. Bender, M. Schreiber, H. Lategahn, T. Strauss, C. Stiller, *et al.*, "Making bertha drive—an autonomous journey on a historic route", *IEEE Intell. Trans. Syst. Magazine*, vol. 6, no. 2, pp. 8–20, Summer 2014, ISSN: 1939-1390.

[20] A. Kesting, M. Treiber, and D. Helbing, "Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity", *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4585–4605, 2010.