Gap Approaching Intelligent Driver Model for Interactive Simulation of Merging Scenarios

Johannes Fischer
Measurement and Control Systems
Karlsruhe Institute of Technology
Karlsruhe, Germany
johannes.fischer@kit.edu

Etienne Bührle
Measurement and Control Systems
Karlsruhe Institute of Technology
Karlsruhe, Germany
etienne.buehrle@kit.edu

Christoph Stiller
Measurement and Control Systems
Karlsruhe Institute of Technology
Karlsruhe, Germany
stiller@kit.edu

***Abstract*—As an important part of automated vehicle development and testing, simulation makes heavy use of driver models to reproduce the behavior of traffic participants. Due to their simplicity, most models fail to capture driver behavior in interactive situations like lane changes or merging, where drivers need to consider multiple vehicles simultaneously and smoothly approach gaps. We propose the Gap APproaching Intelligent Driver Model (GAP-IDM), an extension of IDM that takes an arbitrary number of target vehicles into account and produces realistic behavior for approaching traffic gaps, even when the ego vehicle has to overtake or fall behind target vehicles. To this end, we use a target distance rectification to produce smooth behaviors even for small or negative distances, and to enforce time or distance limits on the maneuver. We evaluate the proposed model in an optional and a necessary lane change scenario and demonstrate that it generates realistic driving behavior. Possible applications of our model include simulations of interactive scenarios, development of complex driver models with multiple target vehicles, or the use as a low-level policy in a high-level behavior planning module.**

***Index Terms*—Intelligent Driver Model, Interactive, Merging, Traffic Gap, Multi-Lane, High-Level Action**

I. Introduction and Related Work

The development of automated vehicles is an important step towards safer and more efficient traffic [1]. Automated vehicles enable novel parking and ride-sharing concepts that can help reduce the total number of vehicles. In addition, they offer the potential to make traffic more efficient by integrating information about other vehicles and their predicted motion, leading to reduced energy consumption [1].

A major part of the development of automated vehicles relies on simulation. For example, planning algorithms are thoroughly tested in closed-loop simulation before being deployed in actual road traffic. This requires realistic models of other traffic participants and their interactions with the ego vehicle.

Simulation of other vehicles can also be part of the planning algorithm itself. For instance, training a reinforcement learning policy or using a Monte Carlo Tree Search-based planning algorithm requires simulating other traffic participants as realistically as possible.

One possibility to simulate the behavior of other traffic participants is to use driver models such as the Intelligent Driver Model (IDM) [2]. Their simplicity and ease of use make them a natural choice for applications that require fast and realistic simulation of driver behavior. Their computational efficiency also makes them interesting for online tree search algorithms or reinforcement learning applications that require large amounts of simulation during training.

(a) Optional lane change due to slower vehicle.

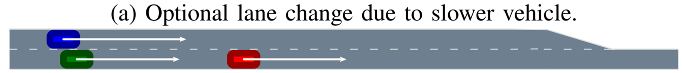

(b) Necessary lane change due to lane ending.

Figure 1: Different situations where the blue ego vehicle approaches the gap between red and green vehicle.

The IDM is a widely used longitudinal car following model that balances driving at a desired speed with maintaining a safe distance to a front vehicle [2]. The model has a small set of interpretable parameters that can express different driving characteristics. For this reason, the IDM has served as a foundation for the development of a multitude of longitudinal driver models, and has been extended by lane change decisions [3], acceleration heuristics [4], risk measures [5], cooperative behavior decisions [6], and multiple targets [7]–[9].

For reasons of interpretability, most driver models are restricted to simple driving situations and are not able to model complex interactions between traffic participants. In particular, we note that currently available, interpretable models fail to capture the behavior of a driver that is approaching a traffic gap to merge into. The ability to approach gaps is an important component of realistic driving behavior, especially in interactive or dense traffic situations that require lane changes, such as highway on-ramps or merging traffic. Furthermore, it forms a basis for interactive, more complex behavior models such as the cooperative IDM [6]. Finally, gaps can be used as high-level action choices for behavior planning [10] and a driver model is a computationally cheap way to simulate the resulting low-level behavior. More generally, the computational cost of low-level rollouts is important for online tree search algorithms or reinforcement learning applications.

To approach a traffic gap successfully, multiple vehicles need to be considered which might not yet be in the right configuration with respect to the ego vehicle. For example, the vehicle at the front end of the gap might not be in front of the ego vehicle at the beginning of the scene, as illustrated in Fig. 1a.

Buyer *et al.* consider a linear combination of target vehicles to determine the ego acceleration [7], [8]. However, this can result in an unsafe driver model: While the original IDM is guaranteed to brake when approaching another vehicle, the linear combination of target vehicles can result in a behavior that collides with a leading vehicle on the ego lane if the target vehicle on a neighboring lane is further away.

This issue can be addressed by considering the maximum over all target vehicles instead of a linear combination [9]. However, none of the models that consider multiple target vehicles [7]–[9] are able to smoothly approach a traffic gap which the ego vehicle is not yet well aligned with, i.e., to which the distance is small or negative. For a small or negative distance, they instead produce a harsh braking maneuver until the ego vehicle is aligned with the gap.

An alternative way to merge into gaps consists of planning the future trajectory to the center of the gap [11]. The disadvantage of this approach is that the behavior cannot be tuned as easily as IDM to reflect different driving styles.

This work proposes the Gap APproaching Intelligent Driver Model (GAP-IDM) which is able to take an arbitrary number of front and rear target vehicles into account, and realistically models driver behavior in cases where the distance to the targets is negative, resulting in smooth gap approaches without harsh braking maneuvers.

We achieve this by introducing a novel rectification technique to handle small or negative distances in IDM. In case of small or negative distances, a rectified value is used in the GAP-IDM formula, which is computed either by a non-negative rectifier function or as the distance to a virtual target vehicle that merges with the real vehicle. This approach enables us to retain a concise parameter set, only adding a comfortable acceleration to the standard IDM parameters, while gaining the ability to smoothly approach traffic gaps.

The presented GAP-IDM can be used in a variety of applications:

- Producing interactive lane changing and merging behavior for traffic simulations.
- Developing more complex cooperative and interactive driver models with multiple relevant vehicles.
- Usage as a high-level action representation for behavior planning with online tree search algorithms or reinforcement learning.

Compared to interactive motion planners [12], [13], GAP-IDM has the advantage of being very fast to evaluate. Furthermore, being based on IDM, it is simple to implement and use and retains the same intuitive parameter set, easily capturing a multitude of driving styles.

Our primary use case in this work is to model the approach towards a traffic gap on a neighboring lane, with which the vehicle is not yet aligned, while maintaining a safe distance to a leading vehicle on the ego lane. While the decision for a gap to be approached is out of the scope of this work, this question can be addressed by a high-level planner that uses the GAP-IDM to simulate the low-level behavior.

We focus on modelling the longitudinal motion for approaching a traffic gap. The lateral motion for merging into the gap is not directly addressed in this work, but GAP-IDM can be easily complemented by lateral motion models such as MOBIL [3].

II. Steady State of the Intelligent Driver Model

The Intelligent Driver Model (IDM) [2] is a microscopic traffic flow model that has gained considerable popularity as a model of longitudinal car following behavior. It satisfies the equation

$$a_{\text{IDM}}(v, s_f, v_f) = a \cdot \left(1 - \left(\frac{v}{v_0}\right)^\delta - \left(\frac{s^*(v, v_f)}{s_f}\right)^2\right)$$

where v is the current velocity, s_f is the distance to the front vehicle, v_f is the velocity of the front vehicle, and the parameter vector $(v_0, s_0, T, a, b, \delta)$ comprises the desired velocity v_0, the minimum distance s_0 to the front vehicle, the desired time headway T, the maximum acceleration a, the comfortable deceleration b, and the exponent δ.

The dynamic desired distance is given by $s^*(v, v_f) = s_0 + \max\left(0, vT + \frac{v(v - v_f)}{2\sqrt{ab}}\right)$. On the one hand, this desired distance maintains the desired time headway in the case of stationary driving ($v = v_f$). On the other hand, the velocity difference term captures an *intelligent braking strategy* that tries to limit braking to the comfortable deceleration [2].

The IDM transitions smoothly between free flow and car following behavior. This leads to the desired velocity also influencing the driving behavior when following a slower vehicle: When following a target vehicle that has a constant velocity v_f, the original IDM converges to the steady state $v = v_f$ and distance

$$s_f = \frac{s_0 + v_f T}{\sqrt{1 - (v_f/v_0)^\delta}}.$$

This is problematic if the desired velocity is close to the target velocity $v_0 \approx v_f$ and can lead to unrealistically large distances in urban driving scenarios [2].

This issue was addressed by separating the free flow and car following regimes in the IDM+ [14], resulting in

$$a_{\text{IDM+}}(v, s, v_f) = a \cdot \min\left(1 - \left(\frac{v}{v_0}\right)^\delta, 1 - \left(\frac{s^*(v, v_f)}{s}\right)^2\right).$$

Since the IDM+ has two separate regimes for free flow and car following dynamics, it has a steady state distance of $s_0 + v_f T$ independent of the desired velocity $v_0 \geq v_f$.

III. Gap Approaching Intelligent Driver Model

The proposed GAP-IDM is an extension of IDM that is capable of considering multiple front and rear targets on different lanes. It is able to smoothly approach gaps in the traffic flow that the ego vehicle is not aligned with at the beginning of the maneuver, as in the case of a front target starting out behind the ego vehicle. Thus, the model realistically captures lane-changing and merging situations where the vehicle has to align itself next to the traffic gap before performing the merge.

This is achieved by first incorporating an interaction term for rear targets in the acceleration equation, similar to the car following term. Next, the equations are generalized to an arbitrary number of target vehicles by considering the vehicle with the maximal interaction term. To address the case where the merging vehicle is misaligned with the gap, we rectify small or negative distances either by using a rectifying function or by introducing a virtual target vehicle.

In the following, we present extensions of both IDM and IDM+ to multiple targets and to approaching misaligned gaps, eventually resulting in the formulation of GAP-IDM.

We denote the free flow term as $\mathcal{F}_{\text{free}}(v) = 1 - \left({}^{v}/_{v_0}\right)^{\delta}$ and abbreviate the desired distance to the front target as $s_f^* = s^*(v, v_f)$. The original IDM can then be stated as

$$a_{\text{IDM}}(v, s_f, v_f) = a \cdot \left(\mathcal{F}_{\text{free}}(v) - \left(\frac{s_f^*}{s_f} \right)^2 \right)$$

while IDM+ is given by

$$a_{\text{IDM+}}(v, s_f, v_f) = a \cdot \min\left(\mathcal{F}_{\text{free}}(v), 1 - \left(\frac{s_f^*}{s_f} \right)^2 \right).$$

We begin the discussion by extending IDM and IDM+ with the possibility to consider rear targets. We proceed by describing how multiple front and rear targets can be taken into account. Finally, we present different methods for dealing with negative target distances, as can occur when merging behind a target vehicle that is not yet in front of the ego vehicle.

A. The Intelligent Driver Model with Rear Targets

The IDM car following behavior is realized by the interaction term $\left({}^{s_f^*}/_{s_f}\right)^2$ which, due to the negative sign, acts as a repulsive force from the front target vehicle on the ego vehicle. We add a similar term for the rear target vehicle that pushes the ego vehicle forward if a rear target approaches.

Let s_r denote the signed distance to the rear target. The sign is chosen such that $s_r > 0$ implies that the rear target is behind the ego vehicle. The desired distance to the rear target is given by $s_r^* = s^*(v_r, v)$. Note that in contrast to s_f^*, the first argument of s^* is the velocity of the target vehicle, thus using the time gap and approaching rate of the rear target vehicle correctly.

As proposed by [9], the IDM can be extended with a rear target by adding the rear interaction term to the equation.

$$a_{\text{GAP-IDM}}(v, s_f, v_f, s_r, v_r) = a \cdot \left(\mathcal{F}_{\text{free}}(v) - \left(\frac{s_f^*}{s_f} \right)^2 + \left(\frac{s_r^*}{s_r} \right)^2 \right)$$

Note that this can lead to the ego vehicle driving faster than its desired velocity. This is in accordance with the behavior of human drivers who also occasionally drive faster than their desired velocity to merge into a faster gap.

Adapting the IDM+ for rear target vehicles requires a slightly different approach. In IDM+ the free flow term is upper bounded by the interaction term $1 - \left({}^{s_f^*}/_{s_f}\right)^2$. To include the possibility for rear target vehicles, we lower bound the free flow term by a corresponding interaction term with respect to the rear target $\left({}^{s_r^*}/_{s_r}\right)^2 - 1$. However, this only works if the rear target interaction term is smaller than the front target interaction term, i.e., if the rear and front target are sufficient far apart. In case they are not, we average the front and rear interaction terms to obtain a continuous acceleration function. This results in collision free behavior with respect to both, the front and rear vehicle, in stationary traffic since an impending collision would lead to $s_f \to 0$ or $s_r \to 0$, resulting in maximum deceleration or acceleration, respectively.

To summarize, the GAP-IDM+ acceleration is given by

$$a_{\text{GAP-IDM+}}(v, s_f, v_f, s_r, v_r) = \begin{cases} a \cdot \max\left(\min\left(\mathcal{F}_{\text{free}}(v), 1 - \left(\frac{s_f^*}{s_f} \right)^2 \right), \left(\frac{s_r^*}{s_r} \right)^2 - 1 \right), & \text{if } \left(\frac{s_r^*}{s_r} \right)^2 - 1 \leq 1 - \left(\frac{s_f^*}{s_f} \right)^2, \\ \frac{a}{2} \cdot \left(\left(\frac{s_r^*}{s_r} \right)^2 - \left(\frac{s_f^*}{s_f} \right)^2 \right), & \text{else.} \end{cases}$$

We now show that GAP-IDM+ is continuous in all arguments.

Proof. We only need to consider the boundary of the case distinction since both defining functions are continuous. To this end, let (v, s_f, v_f, s_r, v_r) satisfy $\left({}^{s_r^*}/_{s_r}\right)^2 - 1 = 1 - \left({}^{s_f^*}/_{s_f}\right)^2$. Then we have

$$a_{\text{GAP-IDM+}}(v, s_f, v_f, s_r, v_r) = a \cdot \left(1 - \left(\frac{s_f^*}{s_f} \right)^2 \right)$$
$$= \frac{a}{2} \left(1 - \left(\frac{s_f^*}{s_f} \right)^2 + \left(\frac{s_r^*}{s_r} \right)^2 - 1 \right) = \frac{a}{2} \left(\left(\frac{s_r^*}{s_r} \right)^2 - \left(\frac{s_f^*}{s_f} \right)^2 \right)$$

□

In particular, both, GAP-IDM and GAP-IDM+, are able to approach gaps that are smaller than required by the desired time gap, which enables to merge into dense traffic.

B. Considering Multiple Front and Rear Targets

In this section, we extend the IDM and IDM+ to take multiple front and rear targets into account. The targets can be on the ego lane as well as other lanes the ego vehicle might change to. For all targets, we consider only longitudinal distances and velocities, independently of the lane.

A very frequent use case is the approach of a traffic gap on a neighboring lane, with which the vehicle is not yet aligned, while maintaining a safe distance to a leading vehicle on the ego lane. In this case, two front targets and one rear target need to be considered.

Wang *et al.* consider multiple front targets by taking the maximum of the interaction terms over all front targets [9]. We use the same idea to also consider multiple rear targets, resulting in the equation

$$a_{\text{GAP-IDM}}(v, s_f, v_f, s_r, v_r) = \\ a \cdot \left(\mathcal{F}_{\text{free}}(v) - \max_{\mathcal{V}_f} \left(\frac{s_f^*}{s_f} \right)^2 + \max_{\mathcal{V}_r} \left(\frac{s_r^*}{s_r} \right)^2 \right)$$

where $\mathcal{V}_f$ and $\mathcal{V}_r$ denote the sets of front and rear targets, respectively.

The GAP-IDM+ acceleration can be adapted in the same way by taking the maximum interaction $\max_{\mathcal{V}} \left({s^*}/{s} \right)^2$ over all front and rear targets, respectively, in the equation.

In case the set of front or rear targets is empty, the maximum is defined as zero for GAP-IDM and as $-\infty$ for GAP-IDM+, essentially removing the corresponding term from the equation.

Due to the maximum operator, at most one front and one rear target have an effect on the resulting acceleration at each point in time. For this reason, we will formulate all equations for a single front and rear target in the remainder of this work. Nevertheless, those equations are to be understood to be valid in the general case of multiple targets as well, by means of the maximum operator, as described above.

C. Approaching Behavior for Negative Distances

A key property of our model is that it produces realistic behavior even when approaching gaps the ego vehicle is not aligned with at the beginning of the maneuver, for instance when a front target starts out behind the ego vehicle. This capability is necessary for being able to approach traffic gaps. Mathematically, such situations are characterized by negative distances to the front or rear targets. In this section, we will show how the proposed model handles these cases.

The issue is addressed by rectifying the distances with a function g such that all interaction terms $\left({s^*}/{s} \right)^2$ are replaced by $\left({s^*}/{g(s)} \right)^2$. The rectifier g must satisfy the following properties:

- $g(s) \approx s$ for all $s > 0$,
- $g(s) > 0$ for all $s \leq 0$,
- g is monotonically increasing.

This ensures firstly that the interaction term does not change for positive distances, and secondly that it is well-defined for negative distances.

We note that introducing such a monotonic rectifier also solves the problem that IDM harshly brakes when the target vehicle changes due to a lane change and the distance to the new target is very small, a problem that was addressed by [15] by introducing a constant acceleration heuristic.

IV. DISTANCE RECTIFIERS

Wang *et al.* use a rectifier $g(s) = \max(s, \varepsilon)$ with $0 < \varepsilon \ll 1$ [9]. This rectifier has the desired properties, but leads to harsh braking with the maximum deceleration as long as the distance to the target vehicle is negative.

We conclude that the properties listed above are only necessary but not sufficient for reasonable driving behavior. Indeed, a more realistic behavior would be for the ego vehicle to smoothly approach the gap within a reasonable time horizon $\tau > 0$. To achieve this, we need to ensure that $g(s) > 0$ is not too small for $s \leq 0$, depending on the desired time horizon τ. In the following, we show different ways of smoothly rectifying target distances in the GAP-IDM and GAP-IDM+ models.

Figure 1 illustrates two types of gap approaches that we distinguish: The first type is an *optional lane change or merge*, with no restriction on the distance at which the lane change has to be completed. In this case, the only requirement is that the gap approach is completed within a reasonable time τ. To fall back behind another vehicle with similar velocity, human drivers will typically not brake, but step off the accelerator, using the engine drag torque to decelerate.

The second type is a *necessary lane change or merge*, which has to be completed within a certain distance due, for instance, to an ending merging lane or an upcoming highway exit. In such situations, human drivers might occasionally exceed the desired velocity to merge into a certain gap in time. Our proposed methods can recover both described human behaviors.

A. Shifted Softplus Rectification

Our first approach is based on the softplus function $\operatorname{softplus}(x) = \log(1 + \exp(x))$, a smooth approximation of $\max(x, 0)$. We introduce a sharpness parameter $\beta > 0$ and shift the argument of the logarithm by $\alpha \geq 0$, such that the rectifier is defined as $g_{\alpha,\beta}(s) = \frac{1}{\beta} \log(1 + \alpha + \exp(\beta s))$. This ensures $g_{\alpha,\beta}(s) > \frac{1}{\beta} \log(1 + \alpha) \geq 0$ for all s and allows to tune the resulting behavior by means of the sharpness and shift parameters. Figure 2 shows $g_{\alpha,\beta}(s)$ and its inverse for different values of α and β.

This rectifier results in smooth driving behavior. However, it is not obvious how to influence the maneuver duration using the parameters α and β. In particular, it is not possible to specify that a gap should be reached after a certain distance. This issue is addressed in the following.

B. Virtual Target Rectification

Our second approach is based on the idea of introducing a virtual target that merges with the actual target vehicle over time. A virtual target is introduced if a new target vehicle is selected on another lane and this new target would lead to a harsh reaction.

More precisely, whenever a new front target vehicle on another lane would lead to an IDM deceleration larger than

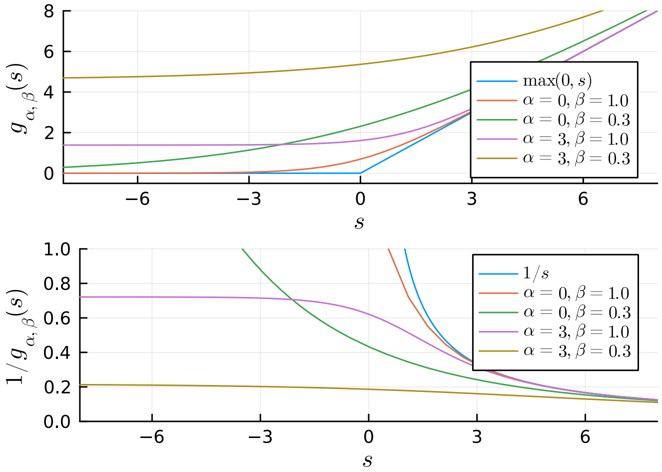

Figure 2: Shifted softplus rectifier $g_{\alpha,\beta}$ and its inverse for different values of α and β.

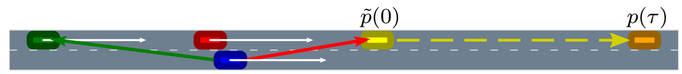

Figure 3: The ego vehicle (blue) wants to merge between the red vehicle and the green vehicle. A virtual target (yellow) is spawned for the front target (red). The virtual target motion (yellow, dashed) is generated from its initial position to the predicted position (orange) of the front target. Until the virtual target merges with the real target, GAP-IDM drives according to the virtual target and the rear target (green), as indicated by the red and green arrows.

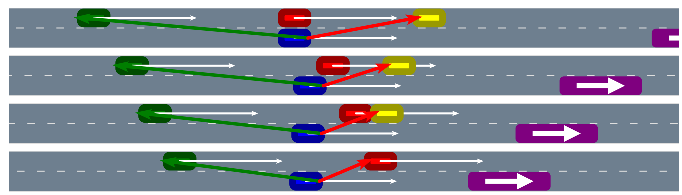

Figure 4: Development of the scene from Fig. 1a at $t = 0\,\mathrm{s}, 4\,\mathrm{s}, 6\,\mathrm{s}, 8\,\mathrm{s}$ using a linear virtual target motion. At $t = 0\,\mathrm{s}$ the virtual target is created until it merges with the real target at $t = 8\,\mathrm{s}$.

the comfortable deceleration b, a virtual target is created, i.e., if $a \cdot \left(1 - \left({s_f^*}/{\max(s_f,0)}\right)^2\right) \leq -b$, or, equivalently, if

$$s_f^* \geq \max(s_f, 0)\sqrt{1 + \frac{b}{a}}.$$

Likewise, a virtual rear target is spawned if a new rear target vehicle satisfies $s_r^* \geq \max(s_r, 0)\sqrt{1 + {c}/{a}}$ where c is the comfortable acceleration.

The GAP-IDM acceleration is then computed by replacing the distance s to the real target vehicle with the distance $\tilde{s}$ to the virtual target vehicle. Thus, the rectifier $g(s) = \tilde{s}$ is defined implicitly as the distance $\tilde{s}$ to the virtual target.

Initially, the virtual target should not have an effect on the ego vehicle. Therefore, it is spawned in a steady driving state with respect to the ego vehicle. As discussed in Section II, the steady state of the regular IDM has issues close to the desired velocity. For this reason, we only apply the virtual target rectifier to the GAP-IDM+ model. The virtual target is initialized with the same velocity $\tilde{v} = v$ as the ego vehicle and a steady state distance of $\tilde{s} = s_0 + \tilde{v}T$.

We generate the motion of the virtual target using one of the procedures described below. At each time step or whenever an acceleration is queried from the GAP-IDM+ model, the motion of the virtual target is recomputed from its planned position, given the new environment configuration. To this end, the planned motion of the virtual target has to be stored between time steps, leading to a stateful rectifier. This allows planning the virtual target motion so as to reach a certain distance after a certain time, by making the maneuver end time part of the stored state and shrinking the horizon τ at each time step accordingly.

To determine the virtual target motion, the real target vehicle is predicted over the time horizon τ using any prediction module. We denote the predicted position and velocity of the real target vehicle with $p(t)$ and $\dot{p}(t)$, respectively. In this work, we use a simple constant velocity prediction. We denote the initial position of the virtual target vehicle by $\tilde{p}(0)$. The procedure is visualized in Fig. 3.

The time horizon τ is an important parameter for adjusting the resulting behavior. In contrast to the stateless rectifier method, we can limit the ego acceleration more precisely and make sure that the gap is reached before the merging lane ends, as described in the following.

In an optional lane change or merge, the time horizon τ can be used to tune the maximum deceleration or acceleration, respectively. In the case of a necessary lane change or merge, the maximum distance until which the lane change has to be completed is given, e.g., by the end of the merging lane. Using the prediction for the target vehicle, this distance can be converted to a time horizon and used for τ. Since the virtual target merges with the real target at the time horizon τ, this parameter can also be used to guarantee that the gap is reached at a desired distance.

1) Linear Virtual Target Motion: One way to generate the motion of the virtual target is to use a linear model. To this end, the future position of the virtual target is linearly interpolated between the initial position $\tilde{p}(0)$ and the predicted target position $p(\tau)$ over the time inverval $[0, \tau]$. At the same time, the velocity of the virtual target is linearly interpolated between the initial velocity $\tilde{v}$ and the predicted target velocity $\dot{p}(\tau)$. An exemplary scene development using the linear virtual target motion is shown in Fig. 4.

2) Jerk-Optimal Virtual Target Motion: An alternative method is to use jerk-optimal trajectories for the virtual target motion. Trajectories given as quintic polynomials are known to minimize the squared jerk [16], resulting in smooth driving behavior of the ego vehicle.

We compute a virtual target motion that interpolates from the initial virtual target position $\tilde{p}(0)$ and velocity $\tilde{v}$ to the predicted target position $p(\tau)$ and velocity $\dot{p}(\tau)$. We fix the initial acceleration to $\ddot{q}_0 = -b$ for front targets and $\ddot{q}_0 = c$ for rear targets and the terminal acceleration to zero.

That is, we compute a quintic polynomial $q(t)$ with initial conditions $q(0) = \tilde{p}(0)$, $\dot{q}(0) = \tilde{v}$, $\ddot{q}(0) = \ddot{q}_0$, and terminal conditions $q(\tau) = p(\tau)$, $\dot{q}(\tau) = \dot{p}(\tau)$, $\ddot{q}(\tau) = 0$.

To meet a maximum acceleration constraint in an optional lane change, polynomial trajectories are computed for a range of τ values. Then their maximum absolute acceleration is computed analytically and the trajectory with the shortest time horizon that satisfies the acceleration constraint is chosen, similarly to [11].

3) Computational Complexity: The addition of virtual targets is not computationally expensive. They are spawned at a fixed position and collision checks are not necessary. Both, the linear and the jerk-optimal virtual target motion, can be computed efficiently [11].

V. Experiments and Evaluation

We evaluate the proposed methods in the AutomotiveSimulator.jl[1] simulation environment. The methods we evaluate are the proposed GAP-IDM approach with the softplus rectifier, the GAP-IDM+ approach with linear and minimum jerk virtual target, and the baseline method of Wang *et al.* [9]. We do not use [7], [8] as baselines since these works are not designed to account for rear target vehicles and are also not safe by design, as illustrated in Section I.

All experiments were carried out using $\alpha = 5, \beta = 0.3$ as parameters of the softplus rectifier, which were tuned empirically. Accelerations are clipped to the interval $[a_{\min}, a_{\max}] = [-9\,\mathrm{m/s^2}, 3\,\mathrm{m/s^2}]$.

A. Optional lane change scenario

We first investigate the GAP-IDM behavior for a randomly created optional lane change scenario in an urban setting, as illustrated in Fig. 1a. To this end, two vehicles are placed on the left lane of a two-lane road and the vehicle that wants to merge into the gap between them is placed on the right lane. We split the evaluation into two settings; the first with the ego vehicle initially placed in the vicinity of the front target, the second with the ego vehicle initially placed in the vicinity of the rear target.

The initial gap between the vehicles on the left lane is sampled from a Gaussian distribution with mean $30\,\mathrm{m}$ and standard deviation $5\,\mathrm{m}$. For the first evaluation, the initial longitudinal position of the merging vehicle is sampled from a Gaussian centered around the longitudinal position of the front target vehicle with a standard deviation of $5\,\mathrm{m}$, such that it can initially be in front of or behind the front target vehicle. For the second evaluation, the position of the merging vehicle is sampled in the same way, but with the Gaussian centered around the rear target vehicle.

[1]https://github.com/sisl/AutomotiveSimulator.jl

The initial velocities of all three vehicles are sampled from a Gaussian distribution with mean $15\,\mathrm{m/s}$ and standard deviation $2\,\mathrm{m/s}$, resembling typical velocities in urban driving.

The two vehicles on the left lane behave according to IDM with the parameters $s_0 = 2\,\mathrm{m}, T = 1\,\mathrm{s}, a = 3\,\mathrm{m/s^2}, b = 2\,\mathrm{m/s^2}, \delta = 4$. The desired velocity of the front target vehicle is sampled from a Gaussian distribution centered around its initial velocity with standard deviation $2\,\mathrm{m/s}$. The desired velocity of the rear target vehicle is set to $v_0 = 18\,\mathrm{m/s}$, to make sure it keeps up with the front target. This results in a nonzero acceleration of the target vehicles and thus in more difficult traffic situations, especially under the constant velocity prediction used in this work. Additionally, the accelerations of the two target vehicles have additive Gaussian noise with standard deviation $0.2\,\mathrm{m/s^2}$. The GAP-IDM vehicle uses the same parameters as the rear target vehicle and $c = 2\,\mathrm{m/s^2}$. The time horizon for the virtual target methods was set to $\tau = 8\,\mathrm{s}$.

We simulate the scene for $20\,\mathrm{s}$ and measure the performance by the mean squared acceleration over the whole trajectory, the total time it takes until the ego vehicle has converged to a steady driving state (defined by $|a|_{\max} \leq 0.15\,\mathrm{m/s^2}$), and the time until the ego vehicle reaches the gap, which is defined as having a minimum positive distance of s_0 to both, front and rear target vehicle.

For each evaluation, we conducted 1000 simulations and report the mean of the performance metrics over all simulations. All four methods were evaluated in the same random initial situations. Exemplary trajectories where all sampled values are set to their mean values, the merging vehicle is next to the front target, and the target vehicles have a constant velocity of $15\,\mathrm{m/s}$ are shown in Fig. 5.

Figure 6 shows the results for the first evaluation, where the merging vehicle is initially placed randomly close to the front target. The baseline method [9] has the highest mean squared acceleration but the lowest time to reach the gap, which is expected since it executes a full braking maneuver whenever it is in front of the front target vehicle. Our proposed methods all have a mean squared acceleration an order of magnitude lower than the baseline method, and still relatively low times to reach the gap. The linear virtual target motion achieves the lowest mean squared acceleration and the softplus rectification has the lowest time to reach the gap. The total time until the merging vehicle has converged to a steady driving state is similar for all methods.

The results for the second evaluation, where the merging vehicle is initially placed randomly close to the rear target, are presented in Fig. 7. Generally, the results are similar to the first evaluation. Due to the acceleration limit of $a_{\max} = 3\,\mathrm{m/s^2}$, the mean squared acceleration of the baseline method is not as high as in the first evaluation, but still highest among all evaluated methods. Of our methods, the linear and minimum jerk virtual targets have the lowest mean squared acceleration, but a longer time to reach the gap.

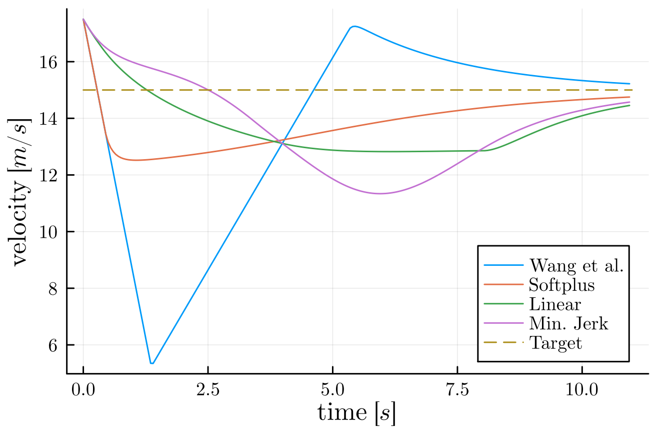

(a) Velocities of the different methods.

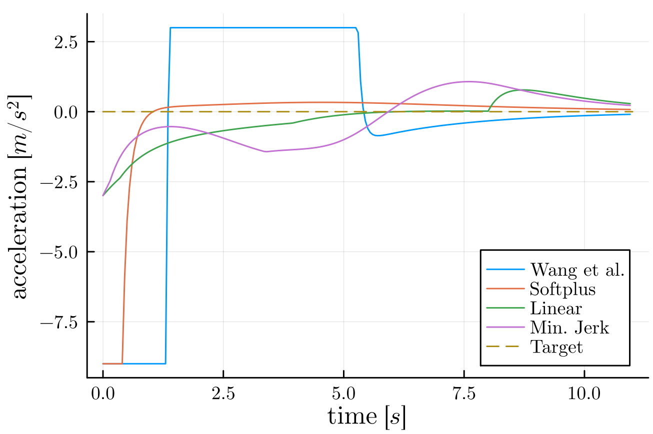

(b) Accelerations of the different methods.

Figure 5: Velocity and acceleration trajectories for the situation in Fig. 1a.

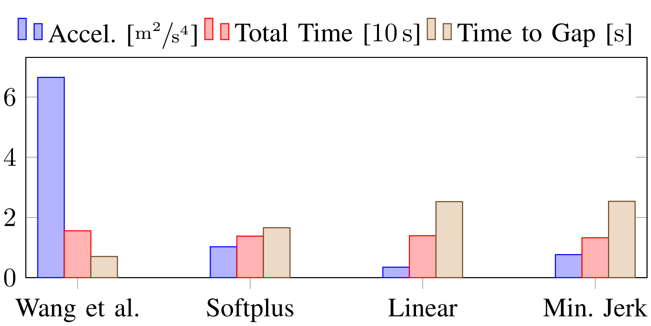

Figure 6: Performance metrics for the optional lane change scenario, where the merging vehicle started close to the front target.

B. Necessary lane change scenario

We evaluate our model in a necessary lane change scenario as typically arises when two lanes merge, such as in Fig. 1b. We randomly sample the distance of the front target to the end of the merging lane from a Gaussian with mean $80\,\mathrm{m}$ and standard deviation $10\,\mathrm{m}$. The remaining parameters (initial

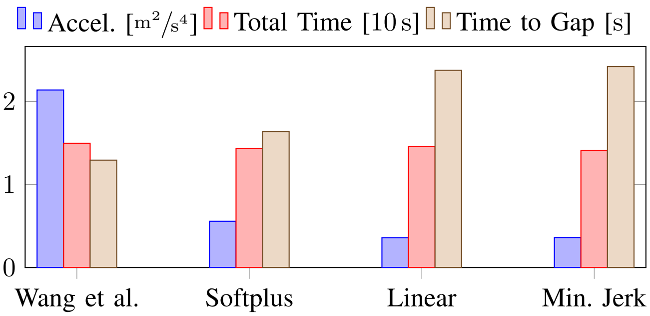

Figure 7: Performance metrics for the optional lane change scenario, where the merging vehicle started close to the rear target.

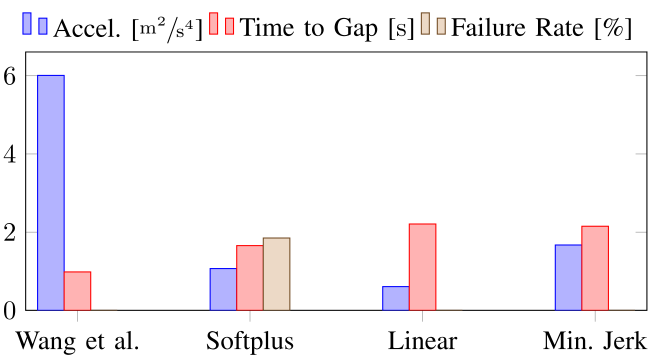

Figure 8: Performance metrics for the necessary lane change scenario.

velocities, distance of the rear target to the front target, position of the merging vehicle, etc.) are sampled as in Section V-A.

We again conduct 1000 simulations with the merging vehicle initially placed near the front target, and 1000 simulations with the merging vehicle initially placed near the rear target. In this scenario we report the failure rate, which is the percentage of simulations where the merging vehicle did not reach the gap before the end of the merging lane.

Figure 8 shows the combined results of all simulations for the necessary lane change scenario. As in the previous experiment, the mean squared acceleration of our methods is significantly lower than in the baseline, with the linear virtual target having the lowest acceleration. The softplus rectification is the only method with a nonzero failure rate of around $2\,\%$.

VI. Conclusions and Future Work

This work introduces GAP-IDM, a simple and computationally cheap driver model for highly interactive traffic situations like lane changes or merging. We extend IDM to handle multiple front and rear targets, taking situations into account where those need to be passed first. This is achieved by rectifying the target distance to a positive value, resulting in smooth driving behavior. Our model can be used to model the gap approaching and merging behavior of a vehicle, even when it is initially misaligned with respect to the targeted traffic gap.

Being based on IDM, our model is collision-free with respect to the front and rear vehicles in stationary traffic.

Our evaluation shows that GAP-IDM is able to approach traffic gaps with a much lower acceleration than the baseline method, resulting in smooth driving behavior that still retains the benefits of IDM. Furthermore, the scenario with a necessary lane change shows that the virtual target rectification approach is able to reach a gap within a certain distance, e.g., when approaching the end of a merging lane. The linear virtual target exhibits lower acceleration values compared to the minimum jerk virtual target, because it moves the target closer to the actual target at a constant rate. In contrast, the minimum jerk virtual target produces a motion with accelerations of larger magnitude in the middle of the time horizon, which in turn requires a stronger reaction of the merging vehicle to stay clear from the virtual target. The stateless softplus rectifier sometimes fails to reach the gap in time, because it is not straightforward to tune its parameters to meet distance constraints. A potential drawback of the virtual target rectification approach is that it requires maintaining an internal state of the virtual target, which we will address in upcoming work.

Another promising future research avenue is to investigate the degree to which GAP-IDM can describe human merging behavior using real driving datasets. In combination with physics-informed neural networks, this allows to learn a policy from real driving data that is regularized with GAP-IDM [17]. Furthermore, the presented model can be used to develop more complex interactive driver models that require to consider multiple vehicles, e.g., to improve the Cooperative IDM [6]. Another exciting future direction is to use GAP-IDM as a low-level policy for behavior planning approaches that use gaps or target vehicles as high-level actions [10].

Acknowledgment

The authors would like to thank the German Research Foundation (DFG) for being funded within the priority program "SPP 1835 – Cooperative Interacting Automobiles".

References

[1] D. J. Fagnant and K. Kockelman, "Preparing a nation for autonomous vehicles: Opportunities, barriers and policy recommendations," *Transportation Research Part A: Policy and Practice*, vol. 77, pp. 167–181, Jul. 2015, ISSN: 0965-8564.

[2] M. Treiber, A. Hennecke, and D. Helbing, "Congested Traffic States in Empirical Observations and Microscopic Simulations," *Physical Review E*, vol. 62, no. 2, pp. 1805–1824, Aug. 2000, ISSN: 1063-651X, 1095-3787.

[3] A. Kesting, M. Treiber, and D. Helbing, "General Lane-Changing Model MOBIL for Car-Following Models," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1999, no. 1, pp. 86–94, Jan. 2007, ISSN: 0361-1981, 2169-4052.

[4] ——, "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.

[5] J. Eggert, F. Damerow, and S. Klingelschmitt, "The Foresighted Driver Model," in *2015 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2015, pp. 322–329.

[6] M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, "Cooperation-Aware Reinforcement Learning for Merging in Dense Traffic," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Oct. 2019, pp. 3441–3447.

[7] J. Buyer, D. Waldenmayer, N. Sußmann, R. Zöllner, and J. M. Zöllner, "Interaction-Aware Approach for Online Parameter Estimation of a Multi-lane Intelligent Driver Model," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Oct. 2019, pp. 3967–3973.

[8] J. Buyer, D. Waldenmayer, R. Zöllner, and J. M. Zöllner, "Data-Driven Merging of Car-Following Models for Interaction-Aware Vehicle Speed Prediction," in *2021 IEEE 24th International Conference on Information Fusion (FUSION)*, Nov. 2021, pp. 1–8.

[9] L. Wang, C. Fernandez, and C. Stiller, "High-Level Decision Making for Automated Highway Driving via Behavior Cloning," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 1, pp. 923–935, Jan. 2023, ISSN: 2379-8904.

[10] B. Mirchevska, M. Hügle, G. Kalweit, M. Werling, and J. Boedecker, "Amortized Q-learning with Model-based Action Proposals for Autonomous Driving on Highways," in *IEEE International Conference on Robotics and Automation*, Xi'an, China: IEEE, 2021, pp. 1028–1035.

[11] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal trajectory generation for dynamic street scenarios in a Frenét Frame," in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 987–993.

[12] C. Burger, T. Schneider, and M. Lauer, "Interaction aware cooperative trajectory planning for lane change maneuvers in dense traffic," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, Sep. 2020, pp. 1–8.

[13] S. Le Cleac'h, M. Schwager, and Z. Manchester, "ALGAMES: A Fast Solver for Constrained Dynamic Games," in *Robotics: Science and Systems XVI*, Robotics: Science and Systems Foundation, Jul. 2020, ISBN: 978-0-9923747-6-1.

[14] W. J. Schakel, B. van Arem, and B. D. Netten, "Effects of Cooperative Adaptive Cruise Control on traffic flow stability," in *13th International IEEE Conference on Intelligent Transportation Systems*, Sep. 2010, pp. 759–764.

[15] M. Treiber and A. Kesting, *Traffic Flow Dynamics: Data, Models and Simulation*. Berlin, Heidelberg:

Springer Berlin Heidelberg, 2013, ISBN: 978-3-642-32459-8 978-3-642-32460-4.

[16] A. Takahashi, T. Hongo, Y. Ninomiya, and G. Sugimoto, "Local Path Planning And Motion Control For Agv In Positioning," in *Proceedings. IEEE/RSJ International Workshop on Intelligent Robots and Systems '. (IROS '89) 'The Autonomous Mobile Robots and Its Applications*, Sep. 1989, pp. 392–397.

[17] Z. Mo, X. Di, and R. Shi, "A Physics-Informed Deep Learning Paradigm for Car-Following Models," *Transportation Research Part C: Emerging Technologies*, vol. 130, p. 103 240, Sep. 2021, ISSN: 0968090X. arXiv: 2012.13376 [cs, eess].