A Tractable Interaction Model for Trajectory Planning in Automated Driving*

J.R. Ziehn[1 2], M. Ruf[1], D. Willersinn[1], B. Rosenhahn[2], J. Beyerer[1 3] and H. Gotzig[4]

***Abstract*— This paper presents an efficient model for combining automotive trajectory planning with predicted environment interactions, named *progressively interacting trajectories* (PITRA). The model allows to plan trajectories for fully-automated vehicles by actively considering how other traffic participants will react to the trajectory, while retaining many of the advantages of variational trajectory optimization methods, in particular expressiveness and ease of computation. This enables maneuvers such as proactively claiming a gap during a lane change in dense traffic, which are impossible to model in classical variational models. The PITRA approach does not rely on a specific prediction model, but can be used in combination with a wide range of existing models. Its model assumptions and limitations are derived theoretically and demonstrated in several realistic scenarios.**

I. Introduction

Several recent approaches (e.g. [1], [2]) have proposed expressing the task of maneuver planning for automated driving through models based on the calculus of variations. In this formulation, a vehicle trajectory ξ is considered theoretically continuous in time and state space (e.g. positions, velocities, steering wheel angles, ...). A *functional* $\mathcal{S}[\xi]$ is defined which maps any possible trajectory onto a real scalar (usually but not necessarily non-negative). The sought optimal trajectory ${}^*\xi$ minimizes $\mathcal{S}$. A key advantage of variational methods is the way they can model desired and undesired properties of trajectories in a unified way, which is intuitive, expressive, computationally tractable and rather well-understood due to its long-established physical relevance.

These models have been advanced in recent publications, to replace purely local trajectory optimization by approximate global optimization (cf. [3], [4]), search space heuristics (cf. [5]) and fail-safe emergency trajectories (cf. [6]). However, a feature that is both in considerable demand (voiced e.g. in the peer reviews to our previous publication, [6]) *and* still absent in the described variational models, is the ability to consider *interaction* during trajectory planning, in the sense that explored trajectory candidates are evaluated with respect to the reactions they cause in other traffic participants. Classical variational models require to define an a-priori environment prediction (possibly fuzzy) that is independent of the ego vehicle's trajectory. In the case of a merging maneuver, such models have to wait for large gaps, instead of (safely) "widening" a narrow gap by initiating the maneuver, expecting other cars to yield (cf. Fig. 1). If they do not, the maneuver can still be aborted, but the ego vehicle is at an advantage by taking into account the effects of its actions. Similarly, the ego vehicle can realize that slowing down before a lane change can cause rear cars to overtake it, thus increasing the risk of the maneuver.

Interaction among traffic participants *without* reference to trajectory optimization has been the subject of diverse research. In these models, a given traffic situation is viewed from the perspective of an "impartial" outside observer, who predicts the future development of the situation, but does not interfere and optimize the trajectory of any involved participant. [7] predicts lane changes by modeling other traffic participants as intelligent agents who base their decision on expected safety distances after the maneuver; [8] predicts traffic situations based on dynamic Bayesian networks and determines the expected behavior of traffic participants based on local "context" parameters such as relative distances and velocities; car following models such as [9] describe how cars adjust their speed with respect to cars ahead; social force models have been applied to the interaction prediction of pedestrians (e.g. [10]) and generalized to include vehicles in [11]–[13]; [14] considers social forces alone as insufficient and therefore combines them with game theoretic models for merging decisions presented in [15]; cellular automata are famously used in the Nagel–Schreckenberg model [16] for freeway traffic and extended e.g. in [17], [18].

In this paper, we propose the *progressively interacting trajectories* (PITRA) model which combines such "impartial" interaction models with the previously described variational models, to incorporate interaction into optimal trajectory planning, while retaining many desirable qualities of the variational models; in particular the efficient computability.

The main focus of this paper is "interaction" in the particular sense of traffic participants knowing only the past trajectories of others, and planning their own future trajectories, as opposed to cases where Car2Car communication is used to *cooperatively* plan maneuvers. The trajectory is planned only for *one* vehicle, the ego vehicle. This however does not rule out the possibility that other vehicles are automated, or communicate state and intentions via Car2Car to the ego vehicle—as long as there is no joint optimization. The methods presented here can be extended to include

*This work has been jointly supported by the State of Baden-Württemberg and the Fraunhofer-Gesellschaft under Grant 7-4332-62-FhG/26.

[1]Fraunhofer IOSB, 76131 Karlsruhe, Germany
{miriam.ruf, jens.ziehn, dieter.willersinn, juergen.beyerer}@iosb.fraunhofer.de

[2]Institut für Informationsverarbeitung (TNT), Leibniz Universität Hannover, 30167 Hanover, Germany
{rosenhahn, ziehn}@tnt.uni-hannover.de

[3]Vision and Fusion Laboratory (IES), Karlsruhe Institute of Technology (KIT), 76131 Karlsruhe, Germany

[4]Valeo Schalter und Sensoren GmbH, 74321 Bietigheim-Bissingen, Germany heinrich.gotzig@valeo.com

PROFILREGION MOBILITÄTSSYSTEME KARLSRUHE

This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft. It is a preprint version of a paper published at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016), and released under the Open Access policy of the Profilregion. The final version of the paper can be obtained from IEEE Xplore Digital Library (ieeexplore.ieee.org).

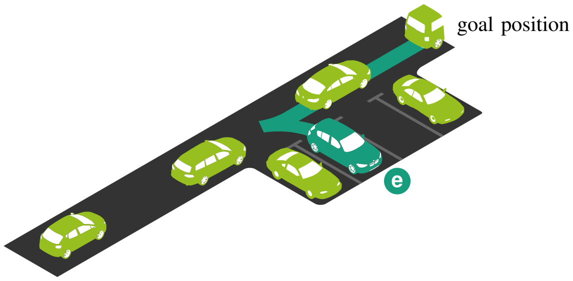

Fig. 1: A motivating scenario for incorporating interaction. The ego vehicle e wants to leave its parking space and merge into flowing traffic, but no predicted gap is wide enough to complete the maneuver before the following car. While classical, non-interactive models would need to wait for a sufficiently wide gap, or use a separate merging logic, an explicit interaction model, as proposed here, could exploit the fact that following cars will likely slow down and yield if the ego vehicle backs up slowly. Two main challenges pertain to this task: Providing models and statistical data for the behavior of other traffic participants, and incorporating these models into an *efficient*, real-time capable trajectory planner. The PITRA model addresses the latter challenge, while allowing for a wide range of possible interaction and environment models.

this case, but by "interaction" we explicitly refer to mutual reactions, not coordinated actions.

Another relevant issue lying beyond the scope of this paper is determining the actual need for interaction models as such. Previous works, such as [1], [2], [19], have considered it sufficient to disregard interaction entirely during a trajectory planning step, and implicitly include it by frequently repeating the planning step after the reactions become apparent. In which cases this approach is sufficient, and where the additional prediction and parametrization effort, that is inevitable in interaction modeling, is outweighed by the improved planning results, cannot be determined without extensive analysis of real-world traffic situations.

This paper focuses on the task of extending variational trajectory planning models to consider interaction with other traffic participants in a way that is efficient to compute, sufficiently expressive and adequate for the safety-critical real-time system that automated driving represents.

II. Problem Description

This section introduces the problem of interactions as addressed in this paper, from the perspective of Euler–Lagrange models (ELMs) which we aim to generalize while retaining their efficiency. Section II-B outlines the classical ELM that does not incorporate interaction; Sec. II-C presents the natural extension of ELMs to incorporate interactions, and concludes, why this model, unlike classical ELMs, cannot be globally optimized efficiently. Based on these considerations, Sec. III can then present an algorithm for reducing the search space to render the solution tractable.

Definition 1 (Time): We consider a real-valued time interval $T = [t_\mathrm{b}, t_\mathrm{e}]$, where t_b is the current time of the maneuver planning algorithm, and t_e is the *planning horizon* of about 5–10 seconds after t_b. Elements of T are denoted t.

Definition 2 (State space): The state space of a dynamic object (a traffic participant or the ego vehicle) is denoted X and describes the relevant properties of the object at each point in time. As states and state changes are part of the model, first derivatives of the properties can be used implicitly; for higher-order derivatives, the state space must be extended. If some states are excluded at certain times, the currently active subset at time t is denoted $X(t) \subseteq X$.

In the models presented here, X is assumed to contain position and positional derivatives (heading, acceleration); depending on the choice of the optimization criteria, the number of dimensions and derivatives can vary. For example, [2] mainly uses a one-dimensional position space, and derivatives up to $\dot{\xi}$ (implicitly: $\ddot{\xi}$), while [1] uses a two-dimensional position space and derivatives up to $\ddot{\xi}$ (implicitly: $\dddot{\xi}$).

Definition 3 (Trajectory): A *trajectory* is a function ξ that maps time onto the state space of a given dynamic object, $\xi : T \to X$, and is assumed to be *sufficiently differentiable* for the required model. The main relevant trajectory in this paper is the trajectory of the ego vehicle, which is assumed to have a fixed starting point $\xi(t_\mathrm{b}) = x_\mathrm{b}$ but a flexible $\xi(t_\mathrm{e})$.

Definition 4 (Subtrajectory): A trajectory ξ can be split along the time axis into *subtrajectories*

$$\xi_a^b : [a, b] \to X \text{ s.t. } \xi_a^b(t) \equiv \xi(t) \text{ for all } t \in [a, b]. \qquad (1)$$

A. Goal Description

The goal is to find a trajectory ξ for the ego vehicle over a given interval T in a given initial situation γ_b that satisfies classical optimality criteria of an Euler–Lagrange model (which will be given in Sec. II-B) and incorporates the reactions of other traffic participants in γ_b on ξ into the planning. The optimization process should be real-time capable with realistic on-board technology.

B. The Non-Interactive Euler–Langrange Model

The classical approach to automated driving is to optimize the trajectory ξ of the ego vehicle, in the sense of "rating" it by a *penalty functional* $\mathcal{S}[\xi] : \Xi \to \mathbb{R}$ and optimizing for ${}^*\xi \in \arg\min_\xi \mathcal{S}[\xi]$. In the worst case, such an approach would require a time complexity of $\mathsf{TIME}\left(|X|^{|T|}\right)$ to optimize globally by evaluating $\mathcal{S}[\xi]$ over all $\xi \in \Xi$, and picking the best-rated trajectory ${}^*\xi$. The computational effort of this brute-force approach is clearly prohibitive. Recent works, such as [1], [2], have chosen a model from the calculus of variations, which expresses the optimality criteria of trajectories by defining a *Lagrangian* L such that the penalty functional of a given trajectory ξ is given by

$$\mathcal{E}[\xi] = \int_{t_\mathrm{b}}^{t_\mathrm{e}} \mathrm{d}t \; L(\xi(t), \dot{\xi}(t), \ddot{\xi}(t), ..., \frac{\mathrm{d}^n \xi(t)}{(\mathrm{d}t)^n}, t). \qquad (2)$$

This functional, called here the *Euler–Lagrange model* (ELM), leads to a trajectory optimization goal which is efficiently soluble by two complementary approaches, an iterative descent from an "initial guess" (cf. [1], [2]), or by discretization and global optimization by transforming (2) into an equivalent Hidden Markov Model (HMM, cf. [3], [6]). The latter provides a complexity of $\mathsf{TIME}\left(|X|^2 \cdot |T|\right)$, which we notice is linear in T. The approaches are compared in [4]. The key to this efficient solution is that the ELM in eq. (2) assumes that all optimality criteria of a trajectory can be established locally along its time parameter. As stated in [1], [2], this is true for a wide range of criteria, including

- comfort and efficiency, by penalizing $|\ddot{\xi}|$ and $|\dddot{\xi}|$,

PROFILREGION MOBILITÄTSSYSTEME KARLSRUHE

This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft. It is a preprint version of a paper published at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016), and released under the Open Access policy of the Profilregion. The final version of the paper can be obtained from IEEE Xplore Digital Library (ieeexplore.ieee.org).

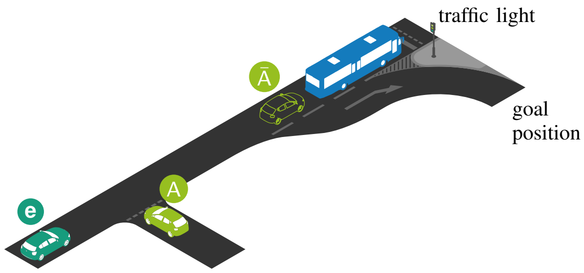

Fig. 2: Potential long-range effects of trajectory decisions. The ego vehicle e has the option of yielding to vehicle A. This decision can adversely affect the future trajectory options: If A lines up at the traffic light behind the bus (candidate position Ā), the ego vehicle must wait as well, since the right turning lane is blocked. If e does not yield to A, it can continue directly to the goal position. An exhaustive search could deliberate this trade-off, but at prohibitive computational effort. State-of-the-art non-interactive models instead are entirely unable to distinguish how the two outcomes depend on the potential behavior of e. The proposed PITRA model provides a compromise by explicitly considering the dependence without exhaustively optimizing the consequences.

- speed limits, by penalizing certain ranges of $\dot{\xi}$,
- reaching a goal, by penalizing certain $\xi(t_\mathrm{e})$, and
- collision avoidance or mitigation with static *and non-interacting* dynamic objects, by penalizing particular tuples of $[\xi, \dot{\xi}, t]$, based on whether other objects are expected to occupy $[\xi, t]$ and at which relative speed.

The assumption of *non-interaction* is vital to the approach, since the locality of the ELM immediately implies that actions cannot have consequences that become relevant later along t. The evaluation requires that the last step along a trajectory can be rated independently of e.g. the first.

C. Extension to Incorporate Interaction

Interaction does not necessarily conform to this assumption. If the ego vehicle decelerates, a rear vehicle may change to the fast lane to overtake it. Even if the ego vehicle accelerates again, the fast lane is now blocked for some time, and the ego vehicle could not change to the fast lane itself without risking a collision. Had the ego vehicle not decelerated earlier, it would be able to change to the fast lane without risk. We conclude that actions of the ego vehicle can, through interaction, have long-term effects on the environment, which in turn affects the trajectory planning (cf. also Fig. 2). In the following, we will formulate and analyze a model to express this effect based on ELMs.

If we consider a state space for the environment Γ with an appropriate Lebesgue measure $\lambda(\gamma), \gamma \in \Gamma$, then the interaction penalty functional can be defined, in analogy to the Euler–Lagrange model, as

$$\mathcal{I}[\xi]=\int_\Gamma \mathrm{d}\lambda(\gamma)\left(p(\gamma|\xi)\cdot\int_{t_\mathrm{b}}^{t_\mathrm{e}}\mathrm{d}t\, L'(\xi(t),...,\tfrac{\mathrm{d}^n\xi(t)}{(\mathrm{d}t)^n},t,\gamma)\right). \quad (3)$$

Here, γ is passed to the extended Lagrangian L', so that for example locations $\xi(t)$ can be penalized based on whether this location is occupied at t in a possible environment γ. All environments are accumulated based on their probabilities given ξ, $p(\gamma|\xi)$. Thus, $\mathcal{I}$ represents the expected penalty given ξ. It should be noted that this formulation presupposes a given initial situation (or situation density), which we will label γ_{t_b} (or $p(\gamma_{t_\mathrm{b}})$), and which corresponds to the observed environment at $t = t_\mathrm{b}$.

Since non-cooperating traffic participants can only base their future actions on the past actions by the ego vehicle (and not its future intentions, like explicit cooperation would permit), and since only the situation at time t, γ_t, is relevant to evaluate L', the model can be simplified to

$$\mathcal{J}[\xi]=\int_{t_\mathrm{b}}^{t_\mathrm{e}}\mathrm{d}t\int_{\Gamma_t}\mathrm{d}\lambda(\gamma_t)\Big(p(\gamma_t|\gamma_{t_\mathrm{b}},\xi_{t_\mathrm{b}}^t,t)\cdot L''(...,t,\gamma_t)\Big). \quad (4)$$

The order of integration is swapped, the entire situation γ is replaced by only the time slice γ_t, whose probability depends only on the portion of the trajectory between t_b and t, labeled $\xi_{t_\mathrm{b}}^t$ (and time t, for formal reasons). Equations (3) and (4) are equivalent based on the assumption of non-cooperating traffic participants, in the sense that $\mathcal{I} \equiv \mathcal{J}$ for all ξ.

It is relevant to note that (4) reduces to an ELM as in (2) if $p(\gamma_t|\gamma_{t_\mathrm{b}},\xi_{t_\mathrm{b}}^t,t) = p(\gamma_t|\gamma_{t_\mathrm{b}},t)$ (i.e. the probability of a future situation does not depend on ξ) by defining

$$L(\xi(t),...,t) = \int_{\Gamma_t}\mathrm{d}\lambda(\gamma_t)\Big(p(\gamma_t|t)\cdot L''(\xi(t),...,t,\gamma_t)\Big), \quad (5)$$

which is identical to the formulation used in [2], [3], [6], [19] (and similar to that used in [1]), a probabilistic yet non-interacting trajectory planning model. Hence, we have established that these previous approaches remain as a special case of the interaction model developed here.

By contrast, neither (3) nor (4) can be reformulated *equivalently* (i.e. without simplifying assumptions) as an ELM over Ξ, so that the efficient solutions presented in [1]–[3], [6] do not apply, and the worst-case effort of $\mathsf{TIME}\left(|X|^{|T|}\right)$ must again be assumed for global optimization.

Another factor not considered so far is the complexity of evaluating the reactions of the environment γ. For a global optimization of (4), the term $\int_{\Gamma_t}\mathrm{d}\lambda(\gamma_t)\ p(\gamma_t|\xi_{t_\mathrm{b}}^t,t)$ must be evaluated for all possible $\xi_{t_\mathrm{b}}^t$, which is $|\xi_{t_\mathrm{b}}^t| \leq |X|^{|T|}$. For each $\xi_{t_\mathrm{b}}^t$, all possible reactions of other traffic participants must be evaluated. An exhaustive search over a set S of traffic participants, each with a state space of $|X_s|$, would require $\mathsf{TIME}((\prod_{s\in S}|X_s|)^{|T|})$, which is generally exponential in both $|S|$ and T. We conclude that no real-time implementation is feasible that evaluates all possible developments this way, and thus, for realistic scales of the traffic situation Γ, S and X_s, the planning horizon T, the state space of the ego vehicle X and the computational power in mass-production automated vehicles, *the global optimization of interacting trajectories is not tractable* with the Euler–Langrange formulations.

III. The PITRA Model

A. Fundamental Considerations

If variational models as the ELM are to be used for trajectory planning, there are mainly two ways to address the issue of computational complexity for interaction modelling:

1) Formulate a simplified functional $\mathcal{S}[\xi]$ that *can* be efficiently optimized globally. This problem would likely not incorporate the full available information and optimality criteria, and thus the globally optimal

PROFILREGION MOBILITÄTSSYSTEME KARLSRUHE

This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft. It is a preprint version of a paper published at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016), and released under the Open Access policy of the Profilregion. The final version of the paper can be obtained from IEEE Xplore Digital Library (ieeexplore.ieee.org).

results would not be truly “optimal” (except within the model). But if its formulation is well-founded, the model limitations and assumptions remain explicit and can easily be validated. One example is the original, non-interacting ELM, but models involving interaction are conceivable as well. The result $\mathcal{S}[^*\xi]$ would generally not be comparable with $\mathcal{J}[^*\xi]$.

2) Use the complete formulation of $\mathcal{J}[\xi]$ as in (4) and introduce a heuristic or local solution method that can efficiently find a solution $^{\circledast}\xi$. Such a solution would be rated directly by $\mathcal{J}[^{\circledast}\xi]$, but not (generally) globally optimize the functional.

The PITRA (progressively interacting trajectories) model explores the second option by extending the Viterbi algorithm which arises from transforming a variational model of automated driving into a Hidden Markov Model (HMM) as presented in [6]. The goal is to establish a sound planning algorithm which is very similar to the classical ELM, yet can consider and exploit expected reactions of other traffic participants. It is rendered tractable by omitting the option to tactically “shape” the state of the environment to best serve a later purpose. Various actions of the ego vehicle are evaluated with corresponding reactions of the environment, but no situation state is intentionally provoked.

B. Discretization and Simplification

The first step is to discretize the model into finite and regularly spaced ego vehicle states $X = \{x_1, ..., x_{|X|}\}$, where each state contains positions and positional derivatives w.r.t. time up to order n, and finite and equidistant time steps $\tau \in \{1, ..., |T|\}$. The trajectory of the ego vehicle $\xi(t)$ thus becomes a sequence ξ_τ. It is shown in [3] that this discretization for an ELM converges to the continuous ELM for infinitesimally fine discretizations, and thus does not constitute a true simplification.

The situation probability distribution $p(\gamma_t|\gamma_{t_\mathrm{b}}, \xi^t_{t_\mathrm{b}}, t)$ is expressed as a situation density obeying the *Markov property*

$$\phi_\tau = \phi(\phi_{\tau-1}, \xi^\tau_{\tau-\underline{\tau}}, \tau) \tag{6}$$

thus assuming it to be uniquely determined by the previous situation density and the recent maneuver of the ego vehicle. The formulation limits the memory of other traffic participants to $\underline{\tau}$ past steps of the ego vehicle. It also limits the memory of other traffic participants concerning their own past states to the number of state properties contained in a situation descriptor γ. It should be noted that the dependence on $\xi^\tau_{\tau-\underline{\tau}}$ means that traffic participants *can* take into account whether, e.g., the ego vehicle changed lanes recently.

Using the definitions of ξ_τ and ϕ_τ, it is possible to redefine

$$\mathrm{d}t\int_{\Gamma_t}\mathrm{d}\lambda(\gamma_t)\Big(p(\gamma_t|...)\cdot L''(...,\gamma_t)\Big)\mapsto\Lambda(\xi_\tau, \xi_{\tau-1}, \phi_\tau, \tau), \tag{7}$$

a Lagrangian that does not rate ξ_τ dependent on derivatives w.r.t. t and a specific situation γ_t anymore, but on the last state transition (of extended, discretized states) and the current situation *density* ϕ_τ. Eventually these definitions replace the penalty functional $\mathcal{J}$ in (4) by a function J

$$\mathcal{J}[\xi] \mapsto J(\xi_1, ..., \xi_{|T|}) = \textstyle\sum_{\tau=1}^{|T|} \Lambda(\xi_\tau, \xi_{\tau-1}, \phi_\tau, \tau) \tag{8}$$

which comprises the discrete evolution and evaluation of the PITRA model as used in the algorithmic solution.

C. Algorithmic Solution

Based on the above diffusion model of the situation, a heuristic can be formulated for planning the trajectory ξ of the ego vehicle. Each subtrajectory ξ^τ_1 implies a situation density ϕ_τ, which in turn can be used to evaluate the next, locally optimal trajectory step $\xi_{\tau+1}$, as shown in Alg. 1. For each time step $\tau \in \{2, ..., |T|\}$ and each state $x \in X$, the optimal preceding state can be chosen. Each state is assigned a ϕ^x_τ, denoting the situation density at time τ in state x given the locally optimal trajectory reaching (x, τ).

Definition 5 (Backtracking functions $\beta, \boldsymbol{\beta}$): The algorithm determines for each time step τ, and each state $y \in X$, an optimal predecessor x (method and optimality will be defined presently), to be stored as $\left[\begin{smallmatrix} x \\ \tau-1 \end{smallmatrix}\right] = \beta(\left[\begin{smallmatrix} y \\ \tau \end{smallmatrix}\right])$. Through this, *backtracking* can be used to establish partial (or complete, for $m = \tau = |T|$) trajectories:

$$\boldsymbol{\beta}(x,\tau,m)=\Big[\underbrace{\beta\circ\cdots\circ\beta}_{m\text{ times }=\,\beta^m}(\left[\begin{smallmatrix} x \\ \tau \end{smallmatrix}\right]), ..., \beta(\left[\begin{smallmatrix} x \\ \tau \end{smallmatrix}\right]), \left[\begin{smallmatrix} x \\ \tau \end{smallmatrix}\right]\Big] \tag{9}$$

Definition 6 (State-dependent situation density ϕ^x_τ): Each state is associated a situation density ϕ^x_τ. This density must *not* be stored for all $\tau \in \{1, ...|T|\}$, but only for the current one (it can be *overwritten*). The index τ thus hints not at the memory size but is kept for clarity.

The algorithm proceeds forward in time from $\tau = 1$ to $\tau = |T| - 1$. At each step τ, for each state $y \in X$, it considers all possible predecessors x, the corresponding ϕ^x_τ, and the transitions $(x, \tau) \to (y, \tau + 1)$, such that the following temporary candidates can be computed:

$$\phi^x_{\mathrm{cnd}} = \phi(\phi^x_{\tau-1}, \big[\boldsymbol{\beta}(x, \tau, \underline{\tau}), y\big], \tau) \tag{10}$$

$$J^x_{\mathrm{cnd}} = \Lambda(y, x, \phi^x_{\mathrm{cnd}}, \tau) + J^x_{\tau-1} \tag{11}$$

Now for each y, an optimal predecessor $x_{\mathrm{opt}} \in \arg\min_x J^x_{\mathrm{cnd}}$ can be found, along with its corresponding situation density $\phi^{x_{\mathrm{opt}}}_{\mathrm{cnd}}$, to set $\beta(\left[\begin{smallmatrix} y \\ \tau+1 \end{smallmatrix}\right]) := \left[\begin{smallmatrix} x_{\mathrm{opt}} \\ \tau \end{smallmatrix}\right]$. As the computation is progressive over τ, all $\beta(\left[\begin{smallmatrix} x \\ \tau \end{smallmatrix}\right])$ used in (10) are defined in time.

The algorithm has a TIME and SPACE complexity of

$$\mathsf{TIME}\left(|T||X|^2(\underline{\tau}+|\phi|)\right) \cap \mathsf{SPACE}(|X|(T + |\phi|)), \tag{12}$$

where $|\phi|$ denotes the size of computing the effect of interaction on the situation density, and depends on the model. If e.g. it uses the full past trajectory of the ego vehicle, then $\underline{\tau} = |T|$, in which case the model would be quadratic in $|T|$. Still, unlike in the complete formulation, it is not exponential.

D. Optimality Considerations

The gain in efficiency with respect to the global optimization of the complete formulation as in (4) implies that information is lost somewhere. Equations (10) and (11) state that any $J^x_{|T|}$ is computed according to (4), which entails that backtracking $\boldsymbol{\beta}(\left[\begin{smallmatrix} x \\ |T| \end{smallmatrix}\right])$ will provide a trajectory which is penalized with respect to the very situations ϕ^x_τ that it is expected to induce through interaction (and all other optimization goals). Therefore the penalty rating is correct, but unlike with the classical Viterbi algorithm for an ELM

This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft. It is a preprint version of a paper published at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016), and released under the Open Access policy of the Profilregion. The final version of the paper can be obtained from IEEE Xplore Digital Library (ieeexplore.ieee.org).

input : state space X, time steps T,
initial situation density ϕ_1^x,
predictor function $\phi(\phi_{\tau-1}, \xi_{\tau-\underline{\tau}}^{\tau}, \tau)$,
Lagrangian $\Lambda(\xi_\tau, \xi_{\tau-1}, \phi_\tau, \tau)$,
start point x_{start}, endpoint type
output: approximately optimal trajectory ${}^{\circledast}\xi$

foreach $x \in X$ **do**
  $J_1^x \leftarrow \begin{cases} 0 & \text{if } x = x_{\text{start}} \\ \infty & \text{else} \end{cases}$;
end

for $\tau \leftarrow 2$ **to** $|T|$ **do**
  foreach $y \in X$ **do**
    foreach $x \in X$ **do**
      // Backtrack the trajectory
      $\boldsymbol{b} \leftarrow \boldsymbol{\beta}(x, \tau, \underline{\tau})$;
      // Update the prediction for $x \to y$
      $\phi_{\text{cnd}} \leftarrow \phi(\phi_{\tau-1}^x, [\boldsymbol{b}, y], \tau)$;
      // Use it to penalize $x \to y$
      $J_{\text{cnd}} \leftarrow \Lambda(y, x, \phi_{\text{cnd}}, \tau) + J_{\tau-1}^x$;
      // Minimize over predecessors x
      if $J_{\text{cnd}}^x < J_{\text{opt}}$ **then**
        $x_{\text{opt}} \leftarrow x$;
        $J_{\text{opt}} \leftarrow J_{\text{cnd}}$;
        $\phi_{\text{opt}} \leftarrow \phi_{\text{cnd}}$;
      end
    end
    // Assign predecessor and corresponding quantities at τ.
    $\beta(\left[\begin{smallmatrix} y \\ \tau \end{smallmatrix}\right]) \leftarrow \left[\begin{smallmatrix} x_{\text{opt}} \\ \tau-1 \end{smallmatrix}\right]$;
    $J_\tau^y \leftarrow J_{\text{opt}}$;
    $\phi_\tau^y \leftarrow \phi_{\text{opt}}$;
  end
end

switch endpoint type **do**
  case *fixed to a given* x_{end}
    ${}^{\circledast}\xi \leftarrow \boldsymbol{\beta}(x_{\text{end}}, \tau, \tau)$;
  end
  case *arbitrary*
    $x_{\text{opt}} \leftarrow \arg\min_x J_{|T|}^x$;
    ${}^{\circledast}\xi \leftarrow \boldsymbol{\beta}(x_{\text{opt}}, \tau, \tau)$;
  end
  case *penalized via a given* $P_{\text{end}}(x)$
    $x_{\text{opt}} \leftarrow \arg\min_x J_{|T|}^x + P_{\text{end}}(x)$;
    ${}^{\circledast}\xi \leftarrow \boldsymbol{\beta}(x_{\text{opt}}, \tau, \tau)$;
  end
end

Alg. 1: Non-optimized PITRA algorithm for trajectory planning. The algorithm evaluates all available states and transitions, assuming that $\Lambda(\xi_\tau, \xi_{\tau-1}, \phi_\tau, \tau) = \infty$ when a given transition or state is not possible at time τ. In dynamic programming, this is easily optimized. As with the classical Viterbi algorithm for variational methods, proposed in [3], the PITRA algorithm features a very high degree of parallelity: The variable y can be parallelized completely, x can be parallelized except for the final minimization. Only the number of time steps (around $|T| = 10$) must be computed serially. For SPACE efficiency, all past situation densities ϕ_τ^x can be discarded. Non-indexed variables exist only within their scope. As with the Viterbi algorithm for HMMs, only the table β of predecessors must be kept for backtracking.

(as in [3]), it is not *exhaustive*, in the sense that the global minimum is not necessarily among any of the backtracking options from any endpoint x in $\left[\begin{smallmatrix} x \\ |T| \end{smallmatrix}\right]$.

Since the algorithm continuously minimizes, the available backtracking trajectories can validly be called *approximations* to the global minimum (calling them *local minima* would be misleading since no notion of locality is used in the algorithm). The quality of this approximation can only be conclusively determined with respect to a given prediction model ϕ.[1] If the prediction model were completely non-interactive (i.e. the estimated actions of other traffic participants can be completely determined *a-priori*), the model would reduce to the classical ELM and the algorithm would yield global optima. The more complex and far-ranging the effects of interaction, the lower the approximation quality, as trajectory candidates are excluded by the following rule:

If two possible trajectories ${}^{a}\xi$, ${}^{b}\xi$ lead to a common point x_1 at the same time t_1 (such that ${}^{a}\xi(t_1) = {}^{b}\xi(t_1) = x_1$), and if $\mathcal{J}[{}^{a}\xi_{t_b}^{t}] < \mathcal{J}[{}^{b}\xi_{t_b}^{t}]$, then only ${}^{a}\xi_{t_b}^{t}$ will be considered further, even if ${}^{b}\xi_{t_b}^{t}$ could produce (through interaction) an environment state ${}^{b}\gamma$ that would be more advantageous in the future (e.g. as shown in Fig. 2).

In principle, assuming a world in which the total rating of a maneuver practically *always* turns out opposite of its initial rating, the result can be arbitrarily bad, since maneuver options are never reconsidered once they are excluded. This is highly unlikely for real-world traffic situations; still, cases such as Fig. 2 can occur where some high-quality solutions are missed. The motivation and general assumption is that on average, the number of considered alternatives and the penalty-minimization mechanisms assure a basic quality of the result. However, as will be stated in Sec. V, an accurate determination of optimality requires a definitive prediction mechanism and substantial real-world traffic data.

E. Comparison with Classical ELMs

The PITRA model is aimed at preserving key advantages that commonly motivate the use of variational models for automated driving. Among these are the expressive and intuitive modeling through Langrangians (as discussed in [1], [2]), the possibility to obtain an efficient solution after a constant number of computations (cf. [3], [6]) and the high degree of parallelism, which invites implementations on GPU or FPGA. Another relevant feature of ELMs, as presented in [3], is that in the case of global optimization, a second, fail-safe emergency trajectory (e.g. to the side of the road) can be computed with almost no additional computational effort. This option is retained in PITRA, where the same principle can be applied, leading to a fail-safe trajectory that also considers interaction. The computational effort is greater due to the need to update predictions several times during planning, instead of just once in the beginning of each

[1] At this point, it should be noted that this consideration relates the PITRA algorithm result to the global optimum of the complete model in (4), not to the *ideal* trajectory in the real-world sense. There is no available benchmark concerning what *ideal* trajectories should look like.

PROFILREGION
MOBILITÄTSSYSTEME
KARLSRUHE

This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft. It is a preprint version of a paper published at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016), and released under the Open Access policy of the Profilregion. The final version of the paper can be obtained from IEEE Xplore Digital Library (ieeexplore.ieee.org).

planning cycle; the complexity with respect to $|X|$ and $|T|$ however can ideally be preserved, or increases just moderately (depending on the complexity of the prediction models). Fuzzy predictions, as used in [2], [6], [19], remain natural in the PITRA model. A challenge lies in parametrizing the interaction models statistically, to realistically estimate the reaction of other traffic participants to the maneuvers of the ego vehicle. Simple a-priori probabilities, as used in classical ELMs, are obtained much more easily.

IV. Practical Application

To show capabilities and limitations of the PITRA model, practical applications are provided in Figs. 3 and 4. In the latter case, a given realistic initial situation is evaluated in PITRA, and several corresponding candidate solutions are shown and discussed in detail. The application uses a prediction method presented in [19] (originally intended for a-priori environment prediction, not for live interaction updating) with parameters set manually to plausible values. This method establishes a Bayesian network among all traffic participants (including the ego vehicle), such that a vehicle a's actions depend on the state of a vehicle b, if b is directly in front of a or on one of the two neighboring lanes (if applicable) or b has the right of way at an intersection or roundabout. This reduces the number of possible interactions to be considered. The result is a positional distribution density for all traffic participants (except the ego vehicle, whose actions are planned, not estimated). For clarity, the fuzzy densities are not depicted in Fig. 4; instead, only the *expected* positions are marked by car symbols. Instances where traffic participants change the state of others by interacting through the Bayesian net are indicated by arrows.

The example shows three candidate solutions (a–d) which are retained until backtracking at the end of the algorithm, since each of them attains a final state that is not better attained by any other evaluated trajectory. Option (e) is discarded early and not pursued further. Which solution, (a), (b) or (c–d), is eventually optimal depends significantly on the penalization of the endpoint ($P_{\text{end}}(x)$ in Alg. 1): If no penalty for lack of progress is specified at all, solution (b) is optimal, because it is perfectly optimal in safety, comfort and fuel efficiency. If the penalty for lack of progress is increased slightly, at some point solution (a) becomes optimal, because it trades off the slight residual risk of a collision, and the discomfort of turning and braking, for significant progress. Only for extreme (and certainly unrealistic) penalties for a lack of progress, solution (c–d) can be obtained, in which case the high collision risk would be balanced by the progress towards the goal. For each case, each trajectory is evaluated and picked based on the expected reactions of the environment, and based on the weights specified just like in the classical ELM. Tradeoffs that are not accepted in the ELM functional are not accepted in the PITRA model either.

V. Conclusion and Outlook

This paper has presented an efficiently computable model for equipping variational trajectory planning methods (Euler–Lagrange models, ELMs) with the ability to anticipate and

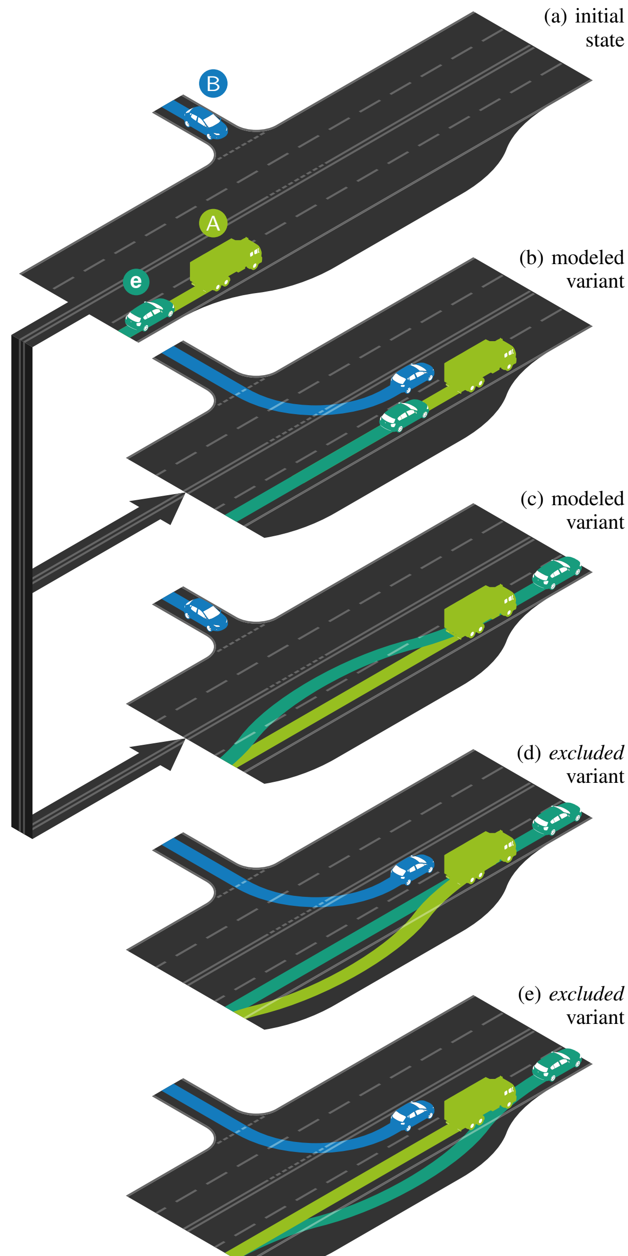

Fig. 3: Model assumptions in PITRA. The ego vehicle e follows a truck A. A car B wants to enter the road via a left turn. Diagram (a) shows the initial state of the situation, diagrams (b–e) show possible developments. As the PITRA model, following the Viterbi algorithm, assumes states at each time τ to uniquely define a past trajectory $\boldsymbol{\beta}(x, \tau, \tau)$, no two trajectories are preserved during planning that lead to the same state at the same time; which trajectories are excluded depends only on their past values of Λ. Therefore, for (b–e) showing the same time τ, (b) and (c) can be retained, but (c–e) are mutually exclusive. The decision that (c) is retained while (d) and (e) are excluded does *not* consider whether B should better enter the road or not. However, case (c) can take into consideration that the maneuver leading e to this place likely keeps B from entering the road.

consider the reaction of other traffic participants, named PITRA (progressively interacting trajectories). It retains key advantages of ELMs, namely their expressiveness, the option for obtaining a definitive trajectory solution after a fixed number of computational steps, which is crucial in safety-critical real-time systems, the high degree of parallelism which enables efficient hardware implementation, and the

PROFILREGION MOBILITÄTSSYSTEME KARLSRUHE

This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft. It is a preprint version of a paper published at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016), and released under the Open Access policy of the Profilregion. The final version of the paper can be obtained from IEEE Xplore Digital Library (ieeexplore.ieee.org).

benefit of providing fail-safe emergency trajectories at almost no added computational effort. The model formulation does not presuppose a specific interaction prediction model; a general expected form for such models is provided.

To develop the concept, a complete but intractable model was presented; based on this model, simplifications in model assumptions and a solution heuristic was presented, whose computation time scales *quadratically* with the ego vehicle state space (as opposed to *exponentially*, in a non-optimized model). Furthermore, the majority of computations can be performed in parallel, leaving a theoretical limit of $|X| \cdot |T| \cdot |\phi|$ sequential steps, where $|X|$ is the size of the ego vehicle state space, $|T|$ is the number of model time steps, and $|\phi|$ is the computational effort of interaction prediction.

The efficiency is achieved by separating the environment state from the ego vehicle state, and exclusively optimizing the ego vehicle state through dynamic programming. This means that the ego vehicle can consider consequences of its actions (such as other vehicles overtaking if it drives too slowly) and exploit them (such as other vehicles yielding to it), but it will not manipulate the environment intentionally to improve its later options, even if this was possible.

The algorithm assures that all evaluated trajectory ratings are *valid*, in the sense that the trajectory would (given the soundness of the prediction models) induce the very situation with respect to which it was rated. The trajectory of the ego vehicle is evaluated this way in each possible state at each predicted time step, and only the progressively optimal $|X|$ candidate trajectories and situations are stored. The ELM is extended to consider all effects of interaction upon the ego vehicle, based on the expressiveness of the applied prediction model. The PITRA model thus limits not the types of interaction, but the search space of ego vehicle trajectories.

The performance of the algorithm was demonstrated on a realistic simulation scenario using a simple prediction model, and on several simple examples. It was shown that the necessary simplifications still allow for a wide range of maneuver options that classically cannot be considered at all.

Outlook

The PITRA model allows to consider interaction in trajectory planning, but whether or not sufficiently expressive interaction models can be parametrized statistically, and whether the effort is worth the improved quality of results, remains an open question. To answer it, comprehensive statistical data sets about traffic behavior must be collected, that allow for a realistic simulation of interaction, and for a stochastically valid parameter estimation.

This model is just an algorithmic framework, whose performance depends significantly on the prediction model that it is combined with. Several such models have been referenced, but their choice was not the focus of this paper. Testing applicable models with PITRA and evaluating their joint performance in realistic scenarios is a critical next step.

An efficient extension to *cooperating* vehicles has yet to be defined; whether such a formulation can and should use the ELM framework, remains to be determined.

REFERENCES

[1] J. Ziegler, P. Bender, T. Dang, and C. Stiller, “Trajectory planning for Bertha – A local, continuous method,” in *Proceedings of the 2014 IEEE Intelligent Vehicles Symposium (IV), Dearborn*, June 2014, pp. 450–457.

[2] M. Ruf, J. Ziehn, B. Rosenhahn, J. Beyerer, D. Willersinn, and H. Gotzig, “Situation Prediction And Reaction Control (SPARC),” in *9. Workshop Fahrerassistenzsysteme – FAS2014*, B. Färber *et al.*, Eds., March 2014, pp. 55–66.

[3] J. Ziehn, M. Ruf, B. Rosenhahn, D. Willersinn, J. Beyerer, and H. Gotzig, “Correspondence between Variational Methods and Hidden Markov Models,” in *Proceedings of the 2015 IEEE Intelligent Vehicles Symposium (IV), Seoul*, Jun. 2015.

[4] M. Ruf, J. Ziehn, D. Willersinn, B. Rosenhahn, J. Beyerer, and H. Gotzig, “Comparison of Local vs. Global Optimization for Trajectory Planning in Automated Driving,” in *10. Workshop Fahrerassistenzsysteme (FAS 2015)*, C. Stiller, K. Bengler, K. Dietmayer, L. Eckstein, B. Färber, M. Maurer, and H. Winner, Eds., Walting im Altmühltal, Germany, September 2015, pp. 71–83.

[5] P. Bender and C. Stiller, “Trajektorienplanung: Manöveridentifikation anhand der Topologie des Freiraums,” in *10. Workshop Fahrerassistenzsysteme (FAS 2015)*, C. Stiller, K. Bengler, K. Dietmayer, L. Eckstein, B. Färber, M. Maurer, and H. Winner, Eds., Walting im Altmühltal, Germany, September 2015, pp. 51–60.

[6] M. Ruf, J. Ziehn, D. Willersinn, B. Rosenhahn, J. Beyerer, and H. Gotzig, “Global Trajectory Optimization on Multilane Roads,” in *18th IEEE International Conference on Intelligent Transportation Systems (ITSC)*, September 2015.

[7] P. Hidas, “Modelling vehicle interactions in microscopic simulation of merging and weaving,” *Transportation Research Part C: Emerging Technologies*, vol. 13, no. 1, pp. 37–62, 2005.

[8] T. Gindele, S. Brechtel, and R. Dillmann, “A probabilistic model for estimating driver behaviors and vehicle trajectories in traffic environments,” in *13th IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2010, pp. 1625–1631.

[9] H. Lenz, C. Wagner, and R. Sollacher, “Multi-anticipative car-following model,” *The European Physical Journal B*, vol. 7, no. 2, pp. 331–335, 1999.

[10] L. Leal-Taixé and B. Rosenhahn, “Pedestrian Interaction in Tracking: The Social Force Model and Global Optimization Methods,” in *Modeling, simulation and visual analysis of crowds*, ser. The International Series in Video Computing, S. Ali, K. Nishino, D. N. Manocha, and M. Shah, Eds. Springer, 2013, vol. 11, pp. 267–294.

[11] W. Huang, M. Fellendorf, and R. Schönauer, “Social Force based Vehicle Model for 2-dimensional Spaces,” in *91st TRB Annual Meeting*, The Transportation Research Board, Ed., 2012, p. 17p.

[12] B. Anvari, M. Bell, A. Sivakumar, and W. Ochieng, “Modelling shared space users via rule-based social force model,” *Transportation Research Part C: Emerging Technologies*, vol. 51, pp. 83–103, 2015.

[13] T. Kretz, J. Lohmiller, and J. Schlaich, “The Social Force Model and its Relation to the Kladek Formula,” in *95st TRB Annual Meeting*, The Transportation Research Board, Ed., 2016.

[14] C. Rudloff, R. Schönauer, and M. Fellendorf, “Comparing Calibrated Shared Space Simulation Model with Real-Life Data,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2390, pp. 44–52, 2013.

[15] H. Kita, “A merging–giveway interaction model of cars in a merging section: a game theoretic analysis,” *Transportation Research Part A: Policy and Practice*, vol. 33, no. 3–4, pp. 305–312, 1999.

[16] K. Nagel and M. Schreckenberg, “A cellular automaton model for freeway traffic,” *J. Phys. I France*, vol. 2, no. 12, pp. 2221–2229, 1992.

[17] L. Lan and C.-W. Chang, “Inhomogeneous cellular automata modeling for mixed traffic with cars and motorcycles,” *Journal of Advanced Transportation*, vol. 39, no. 3, pp. 323–349, 2005.

[18] C. Mallikarjuna and K. Rao, “Cellular automata model for heterogeneous traffic,” *Journal of Advanced Transportation*, vol. 43, no. 3, pp. 321–345, 2009.

[19] M. Ruf, J. Ziehn, D. Willersinn, B. Rosenhahn, J. Beyerer, and H. Gotzig, “A Continuous Approach to Autonomous Driving,” in *Vehicle and Infrastructure Safety Improvement in Adverse Conditions and Night Driving (VISION), Versailles*, Oct. 2014.

PROFILREGION
MOBILITÄTSSYSTEME
KARLSRUHE

This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft. It is a preprint version of a paper published at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016), and released under the Open Access policy of the Profilregion. The final version of the paper can be obtained from IEEE Xplore Digital Library (ieeexplore.ieee.org).

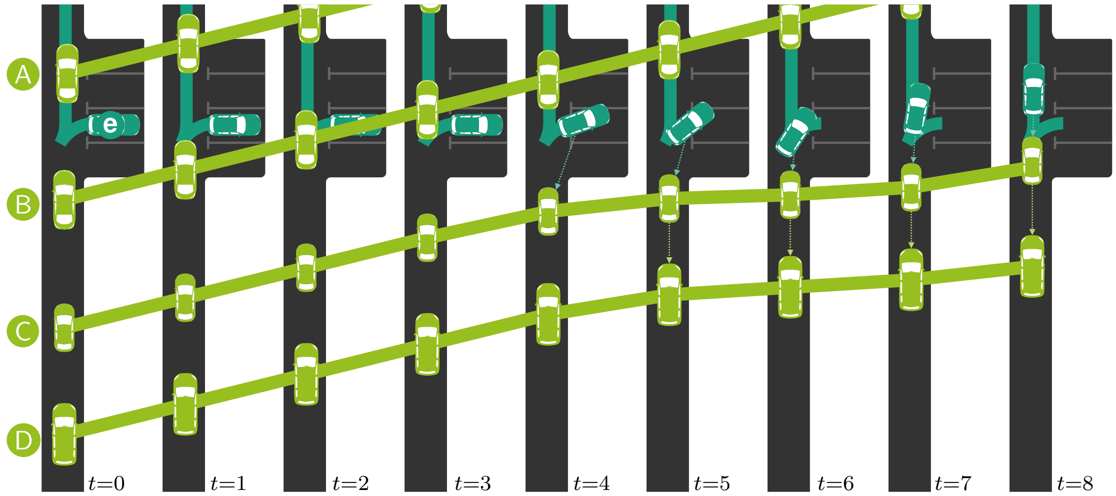

(a) A potential solution trajectory and corresponding expected environment states. The ego vehicle waits for the gap between B and C before backing out of the parking space, causing C to yield. In turn, C braking affects D, who also slows down.

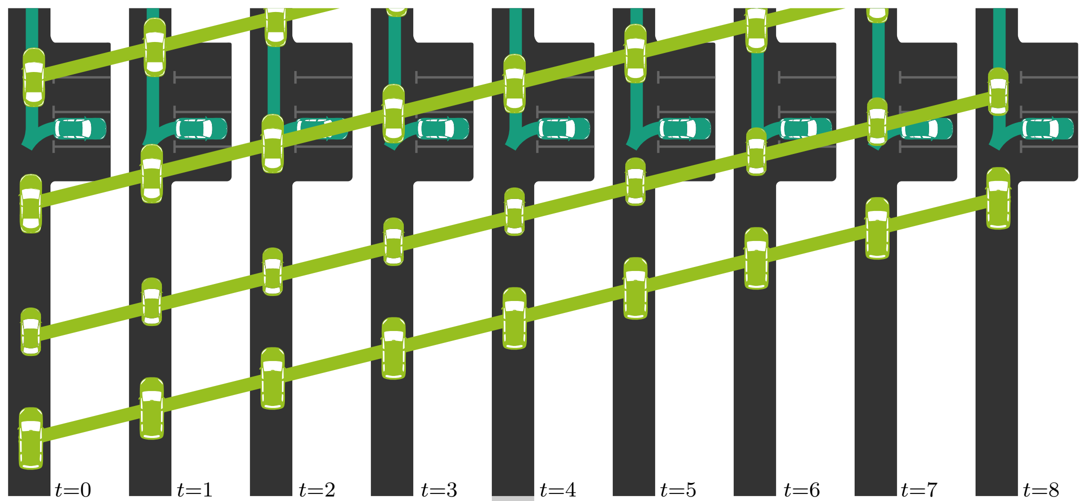

(b) The idle trajectory is optimal in safety and comfort, but is penalized due to a lack of progress. Depending on the weight of this factor, and the penalties necessary to achieve solution (a), either (a) or (b) can be an optimum.

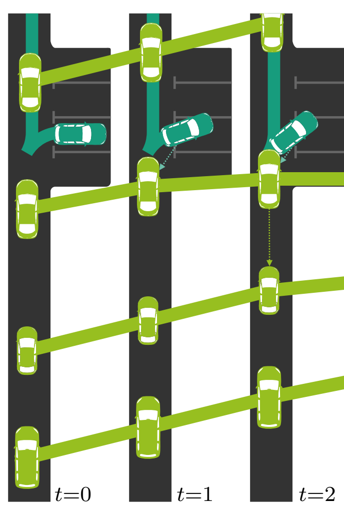

(c) Leaving the space immediately causes a high collision risk with B.

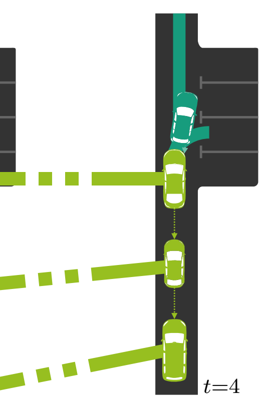

(d) Option (c) is explored further (with t=3 not depicted) as there is no better way to reach this progress at t=4. However, only for very high progress weights, this will be the overall optimal trajectory from which to backtrack.

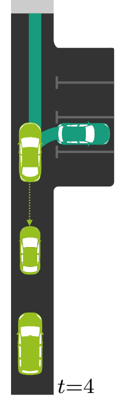

(e) This option of returning to the parking space after (c, t=2) is not explored, because it leads to the same state as (b, t=4), but (b) achieves it at better safety and comfort. Whether this induced a more desirable traffic situation is not considered.

Fig. 4: Practical application to the example in Fig. 1. All times $t > 0$ show the predicted interactions. Positional uncertainties are not shown, car symbols are placed at the *expected position*. Arrows indicate direct interactions causing state changes (e.g. braking). Bold black lines trace positions over time.

PROFILREGION
MOBILITÄTSSYSTEME
KARLSRUHE

This publication was written in the framework of the Profilregion Mobilitätssysteme Karlsruhe, which is funded by the Ministry of Economic Affairs, Labour and Housing in Baden-Württemberg and as a national High Performance Center by the Fraunhofer-Gesellschaft. It is a preprint version of a paper published at the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016), and released under the Open Access policy of the Profilregion. The final version of the paper can be obtained from IEEE Xplore Digital Library (ieeexplore.ieee.org).