A Continuous Approach to Autonomous Driving

M. Ruf[12], J. R. Ziehn[13], D. Willersinn[1], B. Rosenhahn[3], J. Beyerer[12], H. Gotzig[4]

1: Fraunhofer Institute of Optronics, System Technologies and Image Exploitation (IOSB), 76131 Karlsruhe, Germany
2: Vision and Fusion Laboratory (IES), Karlsruhe Institute of Technology (KIT), 76131 Karlsruhe, Germany
3: Institut für Informationsverarbeitung (TNT), Leibniz Universität Hannover, 30167 Hanover, Germany
4: Valeo Schalter und Sensoren GmbH, 74321 Bietigheim-Bissingen, Germany

Abstract: In this paper we discuss the application of the Situation Prediction and Reaction Control (SPARC) concept for fully automated driving, using variational trajectory optimization to avoid static and dynamic obstacles in a real urban traffic situation.
Keywords: Fully automated driving, collision avoidance, trajectory planning

1. Introduction

The Situation Prediction and Reaction Control (SPARC) concept ([Zie12], [RZR+14b], cf. Fig. 2) aims to achieve fully automated driving without the need for extensive discretization of states, both for the ego vehicle and for the enviroment. Some advantages of the concept are:

- The ego vehicle remains in a constant internal state. Discrete state transitions are not necessary, neither are dedicated systems that manage different driving tasks [UAB+08, BBF+08, MBB+08].
- The trajectory planning considers all modeled goals at the same time, including:
  - Safety (collision avoidance)
  - Legality (compliance with traffic rules)
  - Comfort (avoidance of high accelerations)
  - Ecology/economy (avoidance of inefficient maneuvers and wearout)
  - Proximity to goal state

  Thus, the final trajectory tries to satisfy all goals based on their individual weights. In contrast, approaches that divide trajectory planning into several steps (e.g. [MBB+08, KZP+08]) may often override the achievements of the previous step through the modifications of the subsequent one (cf. Fig. 1). The trajectory planning is also fully flexible and not restricted to maneuver primitives (such as *left turn, right turn, lane change, straight ahead*, see [UAB+08] for a more complex example).
- The way in which the predicted development of the situation and other goals are considered together resembles human decision making and should thus lead to fully automated behavior that is still intuitive for human traffic participants.

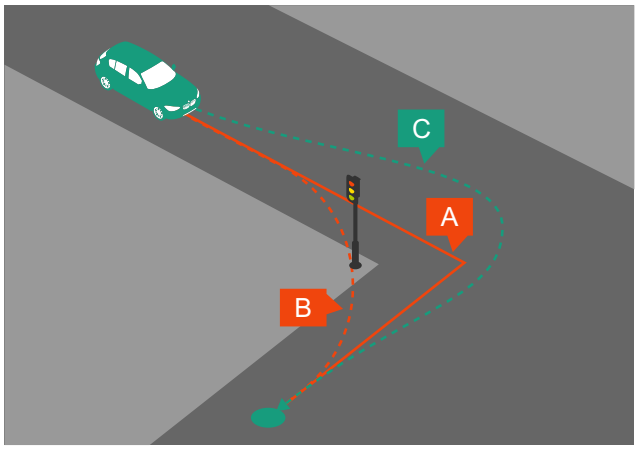

Figure 1: Problems arising from a two-step trajectory computation. **a**: Collision avoidant, but impossible trajectory. **b**: The same trajectory after smoothing—drivable, but not collision avoidant anymore. **c**: A trajectory computed by taking vehicle dynamics and obstacles into account at the same time.

- Partial system failures can be covered by an emergency mode that uses the best available knowledge to bring the ego vehicle to a safe state. Again no separate system is required.

Furthermore the approach *explicitly* trades off between safety goals, such as collision avoidance, and secondary goals, such as comfort and traffic rules. It requires a weighting between the individual goals and thus an explicit assessment of the severity of undesirable states, such as collisions, traffic rule infringements and uncomfortable accelerations, measured on the same scale.
It is important to realize that such a trade-off is a natural part of human driving, and indeed inherent to travel in general (in trading off the risks and gains of locomotion), so it does seem prudent to take this decision explicitly on quantitative grounds for fully automated driving.

2. The Situation Prediction

The situation prediction block uses a dynamic map of the environment to produce a scalar or vector field $\boldsymbol{p}(x, y, t)$ that maps space and time coordinates onto scalar or vector valued *pre-penalties*. These pre-penalties can be turned into actual *penalties* $p \in \mathbb{R}_{\geqslant 0}$ if the ego vehicle

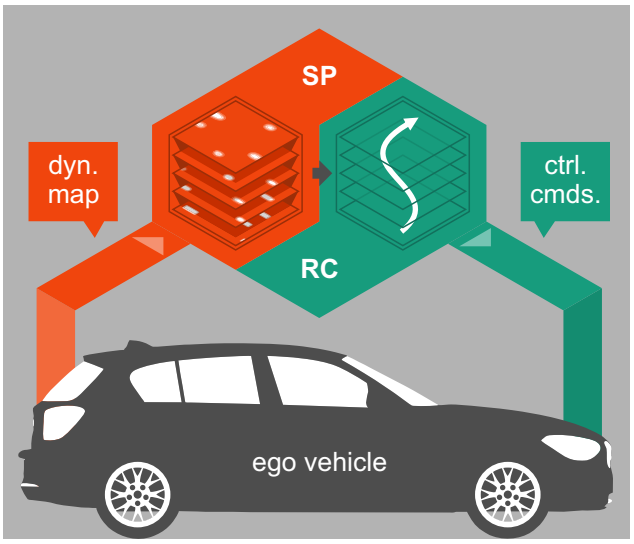

Figure 2: Overview of the SPARC concept as introduced in [RZR+14b]. A dynamic map of the environment (containing both dynamic and static obstacles along with navigation instructions and further priors such as traffic rules) is passed to the *situation prediction* (**SP**) block. The **SP** block computes a *penalty field* $\boldsymbol{p}$ (cf. Sec. 2) and passes it on to the *reaction control* (**RC**) block (cf. Sec. 3), which optimizes the trajectory and passes the very next control commands to the ego vehicle.

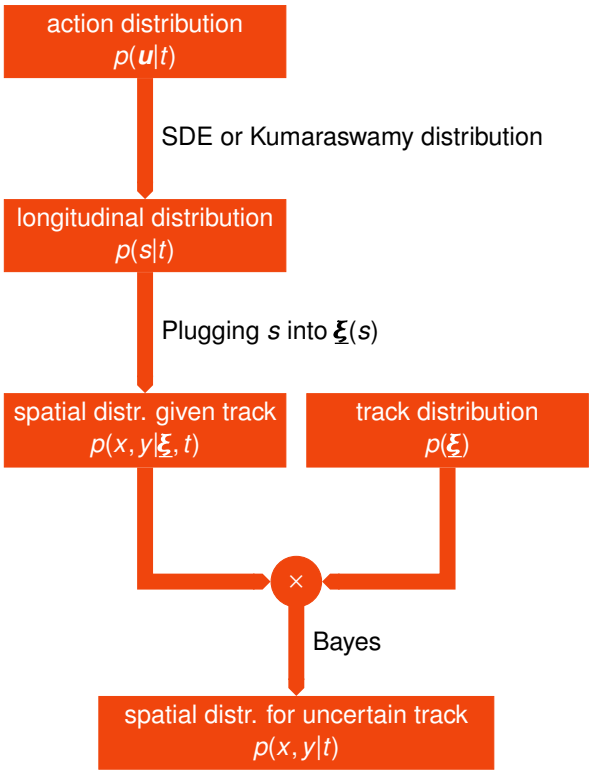

Figure 3: Distributions used to derive occupancy probabilities in the SPARC concept.

plans to pass this particular coordinate along its trajectory $\underline{\boldsymbol{\xi}}(t)$ (see definition at [5]). The actual transformation depends on the choice of $\boldsymbol{p}$. As a simple example using just one other vehicle c, the pre-penalties $\boldsymbol{p}_c(x,y,t)$ may store the probability p of c reaching the position $[x,y]^\top$ at time t, along with the predicted speed $\boldsymbol{v}_c$ that c may attain at $[x,y,t]^\top$. All impossible locations have $p_c = 0$ and undefined $\boldsymbol{v}_c$. For all places traversed by the ego vehicle along a potential trajectory $\underline{\boldsymbol{\xi}}$, the penalty could be computed as follows to be proportional to the expected value of the *kinetic energy* of the impact (refer to [RZR+14b] for a discussion about the use of kinetic energy as a penalty):

$$\boldsymbol{p}_c(x,y,t) = \begin{bmatrix} \boldsymbol{v}_c \\ p_c \end{bmatrix} \quad [1]$$

$$\Rightarrow p_c = p_c \cdot \|\boldsymbol{v}_c - \dot{\underline{\boldsymbol{\xi}}}(t)\|^2 \quad [2]$$

The advantage of using expected values for penalties is that they are additive for independent events and can be weighted by probabilities. The following section will describe how occupancy probabilities can be derived for vehicles.

2.1 From Behavior Distributions to Position Distributions

Vehicle behavior is modeled in several independent steps whose individual methods can be exchanged, provided that the stochastic interpretation is maintained. Figure 3 gives an overview of the distributions that are used in the SPARC concept to determine the probability of an individual vehicle occupying a space/time cell at $[x,y,t]^\top$. The following sections give a detailed account of the individual distributions.

Action Distributions and Bayesian Nets: The first distribution is denoted $p(\boldsymbol{u}|t)$ and describes the probability of the vehicle driver applying certain actions $\boldsymbol{u}$ at time t. Actions can refer to low-level control commands, such as braking, accelerating or steering, or to more abstract maneuvers such as lane changes. There are many possible methods for defining this distribution, depending on the amount of prior information and live sensor data that is available. We propose a Bayesian network over the traffic participants, where the behavior of a vehicle c is modeled depending on vehicles

- that c can observe without using rear-view mirrors and
- that are directly in front of c or on one of the two neighboring lanes (if applicable) or
- that have the right of way at an intersection or roundabout

while for example disregarding the influence of any car that is behind c. This is clearly a simplification, however it allows to reduce the complexity of a traffic situation significantly through a simple set of rules. The advantage is that the behavioral dependencies can be learned in a large variety of training samples, since these topological assumptions are not very restrictive. Furthermore long-distance dependencies will emerge (for example a long

chain of vehicles braking one after another because the foremost vehicle has to stop at a red traffic light), yet only dependencies on direct neighbors must be modeled.
Figure 4 gives an example of such a Bayesian network. D tries to merge into regular traffic. The intended maneuver is clear, only the timing is open. The network models a dependency on the actions of B, but as long as B keeps the distance, D will be assumed to merge (by model assumptions independently of E, since E is behind D). E will react to the merge by either braking (triggering a brake of F) or changing lanes. The latter option is significant for the ego vehicle's behavior. The occupancy probability in front of the ego vehicle increases due to the anticipated behavior of F. The effect of C on D is disregarded because they are not on neighboring lanes. However, long distance effects are modeled such that indirectly, the behavior of A possibly even influences G. The ego vehicle only affects the behavior of G—all other cars are not expected to be aware of the ego vehicle.
In most cases the maneuvers of vehicles are limited to certain discrete *tracks* $\underset{\sim}{\xi}$, which are a function mapping arc length (or any other suitable parameter) to space coordinates,

$$\underset{\sim}{\xi} : s \mapsto [\underset{\sim}{\xi}_x, \underset{\sim}{\xi}_y]^\top \qquad [3]$$

Tracks $\underset{\sim}{\xi}(s)$ (which are just space curves independent of time, such as lanes in a map), relate to trajectories $\underset{\sim}{\xi}(t)$ (which are space/time curves for specific vehicles) via $\underset{\sim}{\xi}(t) := \underset{\sim}{\xi}(s(t))$. If a vehicle is known to be limited to a given track, its *timing* $s(t)$ is still uncertain. To derive this timing from the action distribution, stochastic differential equations [Øks03, All07] would yield sound models, which are however in general difficult to compute.

Longitudinal Distribution: The approach proposed in [RZR+14b] is to approximate speed and position distributions by a Kumaraswamy distribution [Kum80] (similar to the more well-known beta distribution, e.g. [Bul79]) to speed up computation, since the Kumaraswamy distribution CDF can be computed exclusively by additions and multiplications for integral shape parameters a and b:

$$X \sim K(a,b) \Rightarrow p(X < x) = \begin{cases} 0 & x < 0 \\ 1 - (1 - x^a)^b & x \in [0,1] \\ 1 & x > 1 \end{cases} \qquad [4]$$

The shape parameters a and b can be determined to best represent the uncertainty about the future behavior (e.g. increase the probability for braking or accelerating and increasing or decreasing the confidence in this prediction). The Kumaraswamy distribution can then be used to represent the "diffusion" of a vehicle's longitudinal position due to a probability density for accelerations and brakes. The Kumaraswamy distribution is defined on a double-bounded interval, so the locations where the vehicle c is expected to be found in the future is limited (for example by its maximum speed).

Spatial Distribution for a Single Track: The longitudinal probability distribution $p(s|t)$ can directly be turned into a spatial (trajectory) distribution via $\underset{\sim}{\xi}(s(t))$ and by expanding this positional density of the vehicle's *origin* to its full *footprint* (i.e. the bounding box or shape of the vehicle aligned tangentially to $\underset{\sim}{\xi}(s)$). This step may also take into consideration uncertainty about the current position or shape of the vehicle c (e.g. due to limited sensor data).

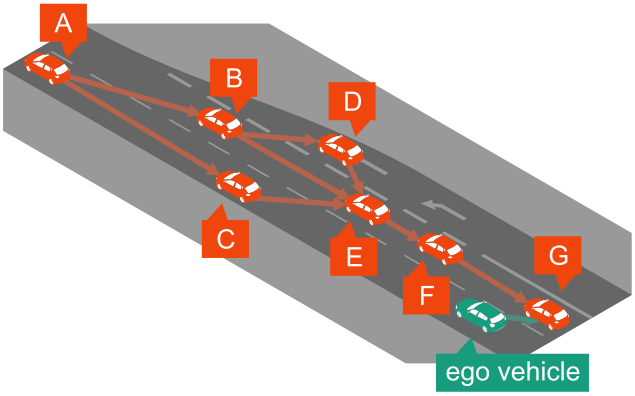

Figure 4: Dependencies as a Bayesian network between vehicles in a merge situation.

Track Distribution for Track Forking: If several tracks $\underset{\sim}{\xi}_c^i$ are possible for vehicle c (e.g. several exits of a roundabout, cf. Fig. 5, referred to as *track forking*), a discrete distribution describing the probabilities of each track is required. This may be a uniform distribution if no prior knowledge exists. Criteria for offline priors (to be assigned to certain intersection exits and, e.g., stored in maps), as exemplified in Fig. 5, include:

- GIS data to determine the relevance of roads from their destinations and intersections
- Navigation map data to determine the relevance of roads from their type
- Actual statistical data gathered at important intersections (for example from traffic surveillance cameras or Car2Car communication)

Live priors that can be observed in an actual scene and be assigned to individual traffic participants would be:

- Turn indicators
- Choice of lanes
- Dynamic behavior, such as slowing down for a highway exit

The combination of both offline and live priors can lead to a distribution $p(i)$ for all $\underset{\sim}{\xi}_c^i$ that are available to c (equivalently written $p(\underset{\sim}{\xi})$ in Fig. 3).

Spatial Distribution for a Track Forking: If the tracks $\underset{\sim}{\xi}_c^i$ are mutually exclusive (i.e. no track $\underset{\sim}{\xi}_c^i$ is fully contained in any track $\underset{\sim}{\xi}_c^j$ for $i \neq j$), the occupancy probabilities can be multiplied $p(i)$ using Bayes' theorem. The resulting probabilities can be used to obtain expected penalties by multiplication and addition as exemplified in [2].

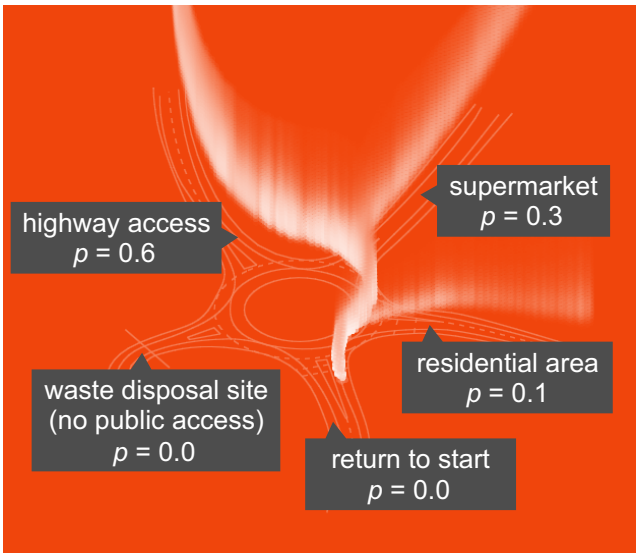

Figure 5: Example of the prediction for a single vehicle about to enter a roundabout (vertical axis is time, same scenario as used in Sec. 4 and Figs. 7 and 8 but from a different perspective). Due to the uncertainty about which exit the vehicle will take, several exits are possible and can be weighted depending on prior information, such as provided by GIS or navigation map analysis. Without proper data, a uniform distribution can always be applied.

2.2 Further Penalty Types

The penalty types discussed thus far only directly relate to collision probabilities with other cars. There are other environmental influences that need to be modeled, although not shown in this paper. These include:

- Models for pedestrian behavior, such as social force models [LTR12], or a simple Irwin–Hall distribution [GKP12].
- Models for traffic rules and road conditions, which can be divided into:
  - Static traffic rules and conditions, such as speed limits, traffic lights, speed bumps, slippery roads, no-passing lines or lanes and legal flow directions in general. These rules can be determined from an offline map and/or sensor input about static objects.
  - Dynamic traffic rules, such as precedence rules, overtaking rules, safety margins. These rules relate to the previously discussed occupancy probabilities, since the future whereabouts of a traffic participant determine whether these rules apply and how.

2.3 General Remarks about the Situation Prediction

The Situation Prediction is a general model to follow the above principles. First of all, do note that it is *not* necessary or suggested to computed penalties for the entire observable space/time volume. On the contrary, only penalties relevant to the trajectory optimization (discussed in the following Section 3) need to be computed. The number of computations depend on the chosen spatial and temporal resolution, and on the optimization method for the Reaction Control. Global methods can be applied to find initial values for iterative, local methods, but for the former, a low resolution evaluation is suggested.
Furthermore, the methods proposed here were chosen for a beneficial trade-off between computational effort and expressiveness. In general the modularity of the approach allows to replace any individual method or any group of methods by different ones. It is suggested that the stochastic interpretation be retained, since a stochastic model allows to make clear statements about the assumptions and limitations of the model. This feature is generally not available for purely heuristic methods, and considered vitally important to assess the applicability of a model and its parameters.

3. The Reaction Control

A trajectory $\boldsymbol{\xi}$ is defined as a space curve that maps time onto $\mathbb{R}^2$ spatial coordinates (cf. [3]):

$$\boldsymbol{\xi} : t \mapsto \begin{bmatrix} \xi_x \\ \xi_y \end{bmatrix} \qquad [5]$$

The trajectory can be interpreted as a curve through space *and* time by considering the point $[\xi_x(t), \xi_y(t), t]^\top$, which will be the usual visualization in this paper. Figure 6 shows a visualization of the Situation Prediction output $\boldsymbol{p}$ along with such a 3D trajectory.
A penalty *functional* $\mathcal{P}$ assigns penalties to trajectories $\boldsymbol{\xi}$ based on $\boldsymbol{p}$, as well as on local trajectory properties that can be expressed in terms of the first few derivatives (in the case of this paper, up to $\ddot{\boldsymbol{\xi}}$):

$$\mathcal{P}[\boldsymbol{\xi}|\boldsymbol{p}] = \int_{t_1}^{t_2} \mathrm{d}t \; p(\boldsymbol{\xi}, \dot{\boldsymbol{\xi}}, \ddot{\boldsymbol{\xi}}|\boldsymbol{p}) \qquad [6]$$

The penalty *function* or *Lagrangian* $p(\boldsymbol{\xi}, \dot{\boldsymbol{\xi}}, \ddot{\boldsymbol{\xi}}|\boldsymbol{p}) \geqslant 0$ evaluates the local penalties based on various properties. The properties used in this paper (also cf. [RZR+14a]) are:

- Collision risks at locations $\boldsymbol{\xi}(t)$, cf. Section 2, referred to as *primary outer penalties* p'_{out}.
- Speeds $\|\dot{\boldsymbol{\xi}}(t)\|$ that exceed the local speed limit or are significantly lower (part of the inner penalties p_{in} that do not depend on the SP block).
- Accelerations $\|\ddot{\boldsymbol{\xi}}(t)\|$ (isotropic in this paper, although for future work a distinction between longitudinal and lateral accelerations is suggested; part of p_{in} as well).

Absent from the discussions in this paper are the *secondary outer penalties* p''_{out} which relate to the Situation Prediction results but not to collision risks (for example localized traffic rules). They are part of the general model but set to zero for the practical applications in this paper.

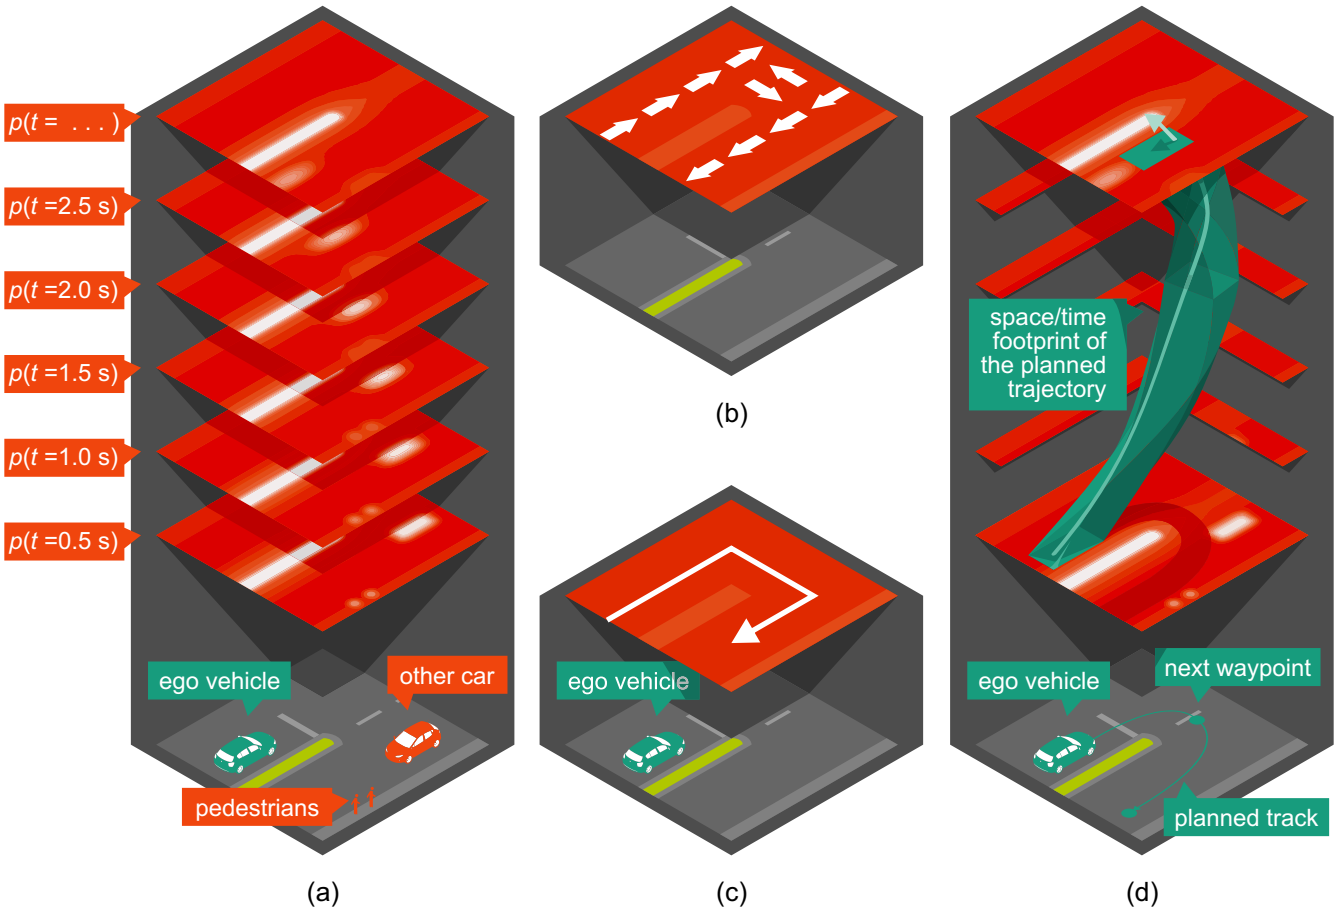

Figure 6: Visualization of the SPARC concept. (a) represents occupancy probabilities (white ⇔ high probability) of objects at several time steps. It can be seen that the other car is predicted to drive ahead, but its timing is "blurred" by the Kumaraswamy distribution. The pedestrians have no known track to follow, so their occupancy probabilities diffuse into all directions. Static objects, as the traffic island, remain the same over all time steps and are blurred only due to sensor uncertainty. (d) represents a trajectory planned by the ego vehicle following given waypoints. The penalties enclosed in the *footprint* (green) along the trajectory are summed and contribute to the total expected penalty of this particular trajectory. (b) and (c) represent traffic rules and navigation instructions that effect the choice of waypoints and influence the total penalty (not part of this paper).

A trajectory $\boldsymbol{\xi}^*$ between two points $\boldsymbol{x}_\mathrm{b} = \boldsymbol{\xi}(t_\mathrm{b})$ and $\boldsymbol{x}_\mathrm{e} = \boldsymbol{\xi}(t_\mathrm{e})$ is called *stationary* with respect to a given functional $\mathcal{P}$ iff

$$\frac{\partial p}{\partial \boldsymbol{\xi}} - \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial p}{\partial \dot{\boldsymbol{\xi}}} + \frac{\mathrm{d}^2}{(\mathrm{d}t)^2}\frac{\partial p}{\partial \ddot{\boldsymbol{\xi}}}\bigg|_{\boldsymbol{\xi}^*} = \delta_{\boldsymbol{\xi}}\mathcal{P}\Big|_{\boldsymbol{\xi}^*} \equiv 0, \qquad [7]$$

where [7] is called the *Euler–Lagrange equation*, that acts as a gradient to optimize the trajectory in a given functional. As this classical definition assumes fixed endpoints, it is as such not suitable for the task at hand, since it is not always desirable to force the arrival at the target position. If the goal position is located such that it can only be achieved at high risk, it should be avoided (such as crossing a busy intersection instead of waiting for a gap, or finding that another vehicle already occupies the target position itself). The following additional terms, known as *natural boundary conditions*, allow a flexible endpoint and assure, that the solution $\boldsymbol{\xi}^*$ is stationary also with respect to its position.

$$\frac{\partial p}{\partial \dot{\boldsymbol{\xi}}}\bigg|_{t_\mathrm{e}} = 0 \qquad \frac{\partial p}{\partial \ddot{\boldsymbol{\xi}}}\bigg|_{t_\mathrm{e}} = 0 \qquad \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial p}{\partial \ddot{\boldsymbol{\xi}}}\bigg|_{t_\mathrm{e}} = 0 \qquad [8]$$

Flexible endpoints introduce the need for an additional penalty inside p, namely p_target, which penalizes deviations from the target point—if such a penalty were not introduced, the safest solution would (usually) be not to move the car at all.

Beside this generalization, the Euler–Lagrange solution must also be constrained to assure that the obtained trajectory can be achieved by the vehicle. The following constraints on steering angles δ, velocities v and accelerations a are proposed (see [RZR+14a] for a derivation of vehicle parameters from given trajectories):

$$\delta = \arctan(\kappa \cdot b) \overset{!}{\in} [-\delta_\mathrm{max}, +\delta_\mathrm{max}] \qquad [9]$$

(where $\kappa = \det[\ddot{\boldsymbol{\xi}}, \dot{\boldsymbol{\xi}}] \cdot \|\dot{\boldsymbol{\xi}}\|^{-3}$ is the local curvature and b is the wheel base of the car)

$$v = \frac{\left\langle \dot{\boldsymbol{\xi}} \middle| \dot{\boldsymbol{\xi}} \right\rangle}{\|\dot{\boldsymbol{\xi}}\|} = \|\dot{\boldsymbol{\xi}}\| \overset{!}{\in} [0, v_{\max}] \quad [10]$$

$$a = \frac{\left\langle \dot{\boldsymbol{\xi}} \middle| \ddot{\boldsymbol{\xi}} \right\rangle}{\|\dot{\boldsymbol{\xi}}\|} = \dot{v} \overset{!}{\in} [a_{\min}, a_{\max}] \quad [11]$$

3.1 Structure of the Penalty Function

The penalty function p in this paper is structured in the following way:

$$\begin{aligned} p(\boldsymbol{\xi}, \dot{\boldsymbol{\xi}}, \ddot{\boldsymbol{\xi}} | \boldsymbol{p}, \boldsymbol{x}_e) = {} & p_{in}(\boldsymbol{\xi}, \dot{\boldsymbol{\xi}}, \ddot{\boldsymbol{\xi}}) \\ & + p'_{out}(\boldsymbol{\xi} | \boldsymbol{p}) \\ & + p''_{out}(\boldsymbol{\xi} | \boldsymbol{p}) \\ & + p_{target}(\boldsymbol{\xi}(t) | \boldsymbol{x}_e) \cdot \hat{\delta}(t - t_e) \end{aligned} \quad [12]$$

where $\hat{\delta}(t)$ denotes the Dirac impulse function, used to switch on p_{target} exclusively at t_e such that its value is added to $\mathcal{P}$. It should be noted that the Dirac delta transforms into a well-defined, finite value through the discretization of $\boldsymbol{\xi}(t)$ over time in the optimization.
The solution thus has the following properties:

- It starts from a given point $\boldsymbol{x}_b$ (hard constraint)
- It leads towards the vicinity of an intended end point $\boldsymbol{x}_e$ (soft goal)
- It is stationary with respect to the penalty functional $\mathcal{P}$ (in particular considering the pre-penalties $\boldsymbol{p}$, hard constraint)
- It exclusively uses steering angles, velocities and accelerations from given, physically possible intervals (hard constraint)

A solution to satisfy the above criteria can be obtained using sequential linear or quadratic programming [NW06, GK02] (SLP, SQP) since both $\mathcal{P}$ and the constraints [9], [10] and [11] are twice continuously differentiable (assuming p'_{out} to be numerically differentiable). A remaining problem is the convergence towards purely local minima; a well-chosen set of several initial trajectories is assumed to resolve this problem for practical applications, however no conclusive evaluation has been performed. The application in Section 4 starts out from a single initial trajectory that is the constant speed connection between $\boldsymbol{x}_b$ and $\boldsymbol{x}_e$. For all timesteps $t > 0$ the previously optimized trajectory is reused.

4. Application

The algorithm is applied to a real-world roundabout scenario (see Figs. 7 and 8 along with these explanations) to demonstrate the process and discuss the individual steps on a practical example.

4.1 Setup, Data and Parameters

A video of approximately 10 seconds was recorded using a GoPro Hero 1 attached to the windscreen of the ego vehicle. There are three cars visible in the scene, labeled **I**, **II** and **III** (according to the order in which they entered the roundabout). In the frames, the car positions were manually detected and projected onto a detailed map of the scene. The prediction of occupancy probabilities was performed by using the current speed of the car as the mode for the Kumaraswamy distribution, and the forking at each exit was computed based on the prior distribution demonstrated in Fig. 5. Figure 7 shows the full predicted occupancy probabilities of the scene up to 20 seconds into the future, to convey an intuition for their distribution over space and time. Figure 8 shows an extract of the same occupancy probabilities along the given track $\underline{\boldsymbol{\xi}}$ for the ego vehicle. Only these occupancy probabilities are relevant for the computation of the optimal trajectory; all others need not be computed explicitly. Furthermore, Fig. 8 shows the trajectory that is planned along this track by the ego vehicle at every time step.

4.2 Discussion of the Individual Steps

Due to the fact that the ego vehicle started at rest and had 20 seconds as an initial time to reach the goal $\boldsymbol{x}_e$, the algorithm converged to a solution that waited for all the traffic participants to pass before entering the roundabout. Had the ego vehicle started at 50 km/h, the obtained solution would enter the roundabout immediately. From its position, vehicle **I** is predicted with certainty to leave the roundabout at the next exit. Vehicle **II** has several possible exits to choose from, so does **III**, who has not reached the roundabout yet. Track forking is weighted in accordance with Fig. 5. The assumption that a vehicle would not return to its entry point (or perform several circles in the roundabout) is merely introduced to better demonstrate the principles; it is not required in any way, nor is it actually suggested, since it may be dangerous in practice.
At $t = 4$, vehicle **II** is observed to leave the roundabout; Figure 7 clearly shows that the left fork of the prediction disappears and instead all the weight is assigned to the exit fork.
Similarly, at $t = 5$, vehicle **III** passes an exit and the weight is assigned to the two remaining exits. This leads to an increased collision risk directly in front of the ego vehicle, since **III** is more likely to pass before it.
At $t = 7$, **III** passes the next exit and is now known to pass before the ego vehicle, again increasing the collision risk ahead of it, and removing the last fork from its prediction. The remaining uncertainty for the occupancy of **III** exclusively relates to its longitudinal behavior (modeled by the Kumaraswamy distribution, leading to the increased blur over time) and the sensor accuracy (leading to a slightly unsharp outline of the "cloud" in space.
After $t = 9$, when **III** is past the ego vehicles ramp, the ego vehicle starts to follow it at roughly the same pace. It should be noted that the planned distance from the most

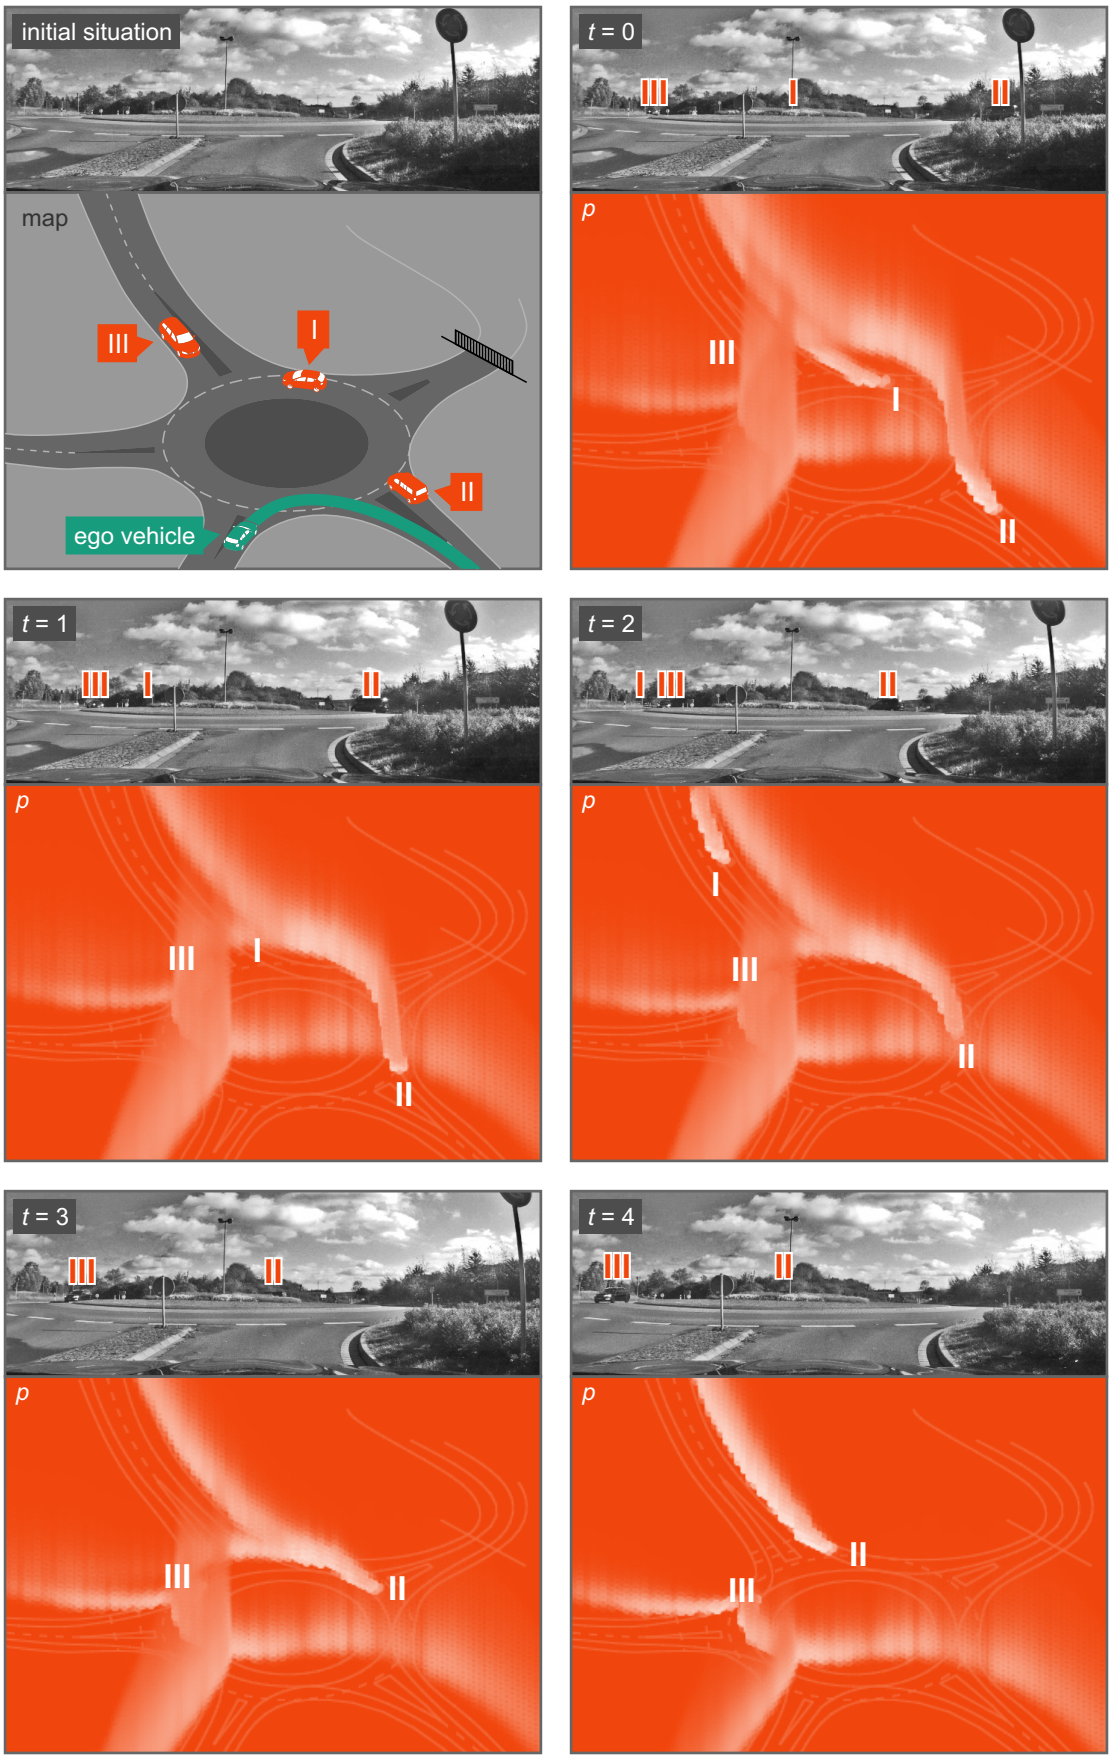

Figure 7: The first image shows the empty roundabout, along with a map of the scene stating the vehicle designations used in the explanations, and the track $\underline{\boldsymbol{\xi}}$ for the ego vehicle (green). The other images show the summed occupancy probabilities p stored in $\boldsymbol{p}$ of the scene at each time step, along with the corresponding dash cam stills for reference. (Cf. Fig. 8.)

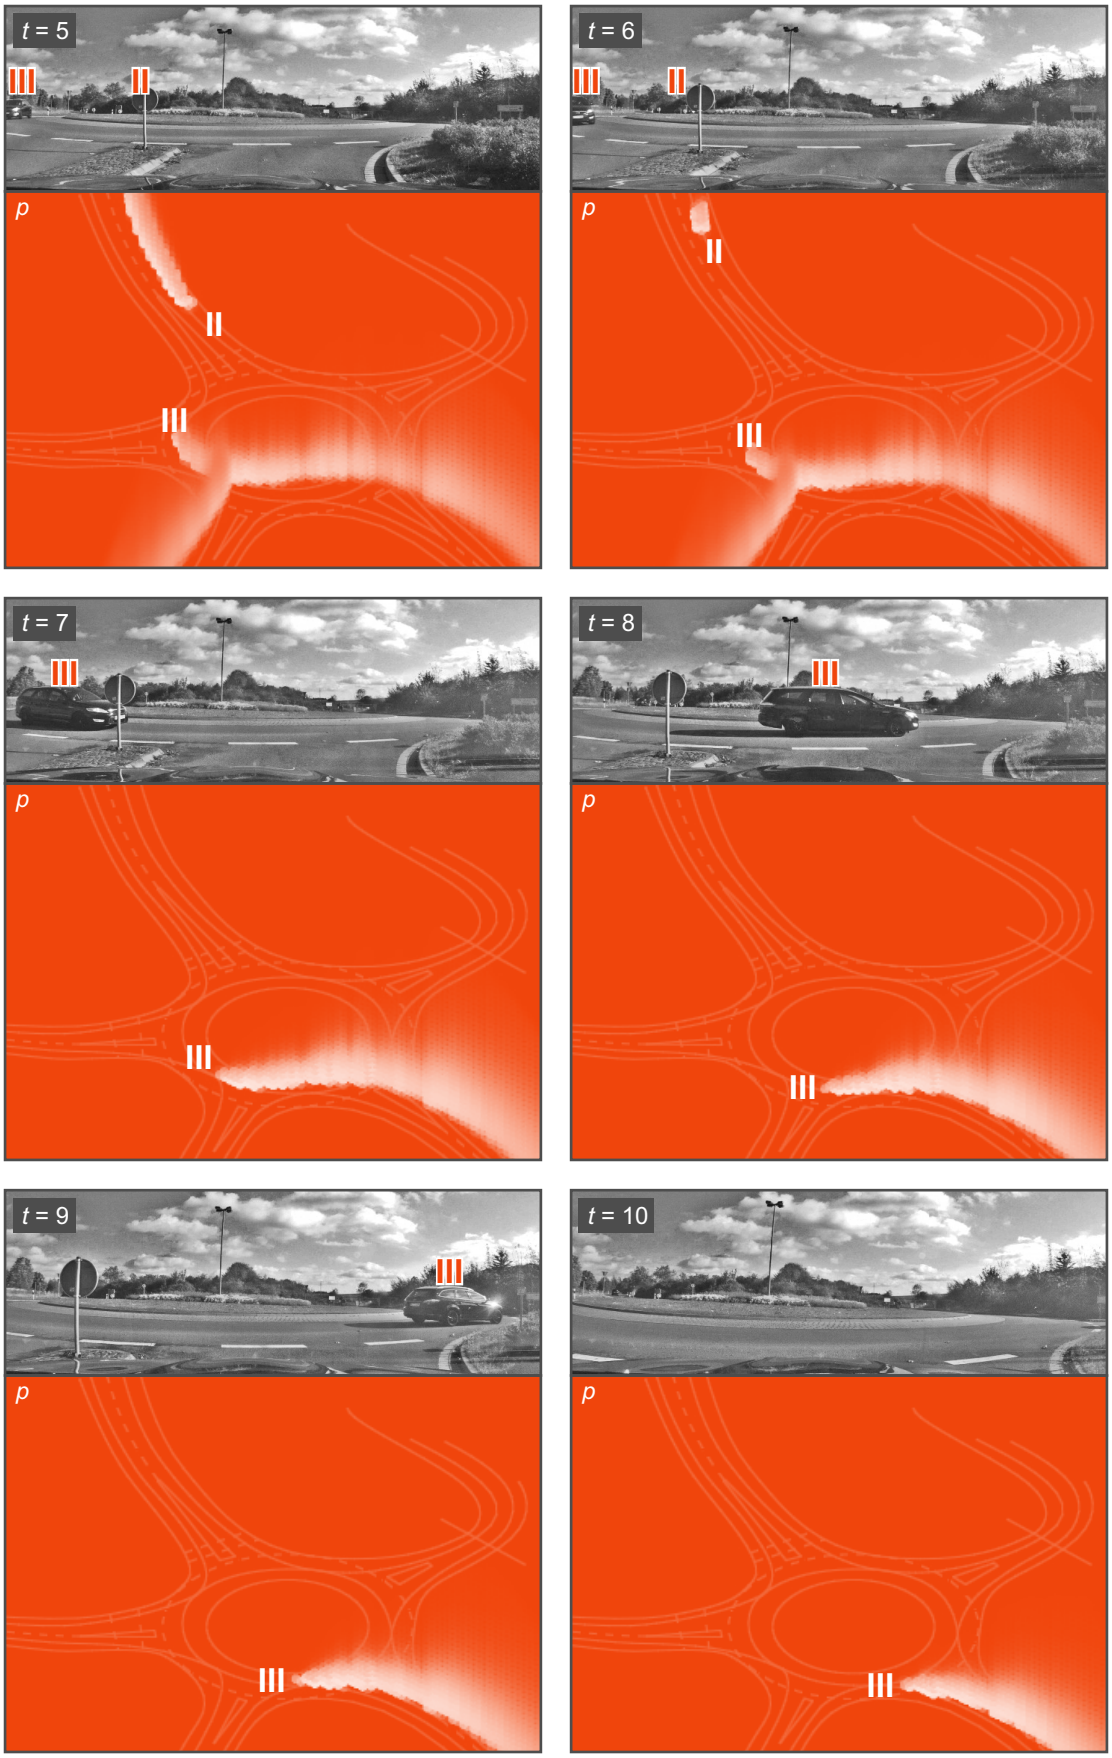

Figure 7: *(Continued from previous page.)*

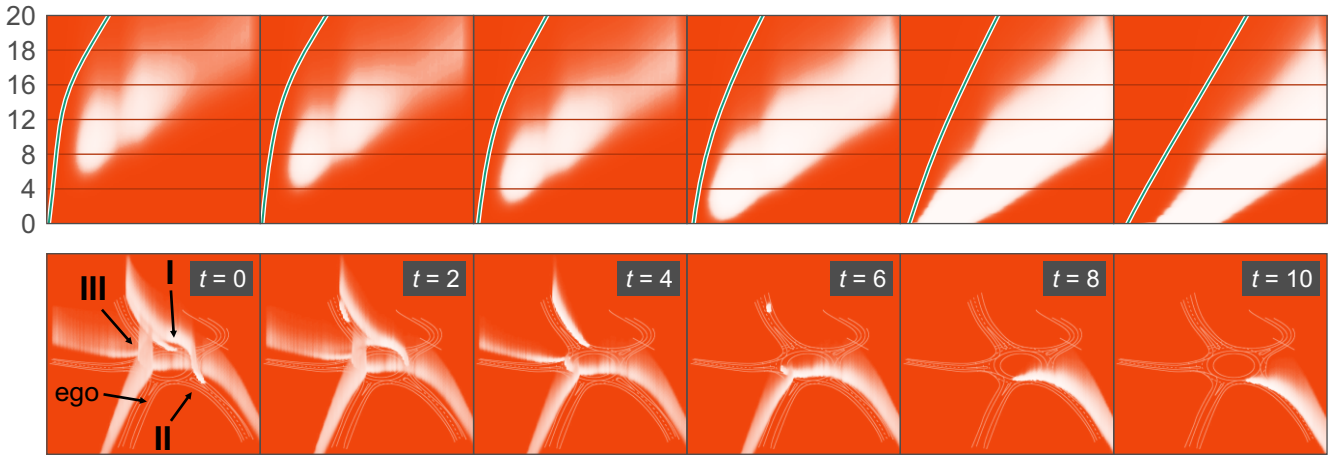

Figure 8: Extracts from $\boldsymbol{p}$ along the given track for the ego vehicle. The top row shows the extracts of p cut out along the ego vehicle's given track along with the optimized trajectory (green). The horizontal axis is track progress; the vertical axis shows predictions in seconds at a finer-than-usual sampling for demonstration purposes. The predicted occupancy probabilities can be seen to become increasingly blurred at more distant prediction times. The bottom row shows corresponding overviews of p.

likely predicted position of **III** (the mode of the "cloud") increases because the uncertainty about the position increases. The algorithm thus plans for the case the **III** might proceed slower than predicted. Should this not be the case, future planning steps would increase the speed to keep a constant distance at all times.

Current Optimum and Fail-Safe Solution: This indicates a particular feature of the SPARC concept, namely that the planned trajectory is supposed to be optimal, given the modeled goals, the model assumptions and the *current* knowledge, not strictly requiring that re-planning will be possible. In case of a partial system failure, a trajectory planned in this way according to the current knowledge should be the optimal way to "spend" or bridge the next (e.g.) 20 seconds. The only aspect that would likely not be desirable is heading for $\boldsymbol{x}_\mathrm{e}$, unless the system failure would be known to be resolved before the planned time horizon is used up. The reason is that $\boldsymbol{x}_\mathrm{e}$ likely lies amid flowing traffic and is thus an unsafe location to reach in a disabled state.
It is suggested to always identify a safe and close location $\boldsymbol{x}_\mathrm{emergency}$ to bring the car to a stop (for example the road shoulder) and plan an emergency trajectory to reach this point, along with the required control commands to perform this trajectory (refer to [RZR+14a] for a preliminary assessment of offline trajectory planning). Should all sensors, the Situation Prediction and/or the Reaction Control fail, this backup set of control commands could be automatically passed to the actuators to reach $\boldsymbol{x}_\mathrm{emergency}$ safely before the predicted time interval runs out.

5. Conclusions and Outlook

As with [RZR+14b], the application of the SPARC concept to individual scenarios yields promising results both in terms of the obtained trajectories for the ego vehicle and for the computational effort. While as of now, no strict optimization of the underlying algorithms was attempted, it is clear that the approach presents a high degree of parallelism and a major part of the mathematical models reduces to multiplications, sums and combined multiply-accumulate operations, for which efficient hardware implementations exist.

The main challenges currently represent the lack of real world data to provide parameters for the applied probability distributions, to test the applicability of models to actual traffic scenarios, and to automatically evaluate the approach using a large database of real-world scenarios. The creation of such a database will be a relevant step to validate the approach and thus represents urgent future work.

Similarly, the results obtained in [RZR+14a] have been used at several places in this work, even though they also currently await validation using real-world measured data. This applies in particular to the emergency trajectory planning indicated in the previous section.

It can be concluded that the SPARC concept shows potential to cope with complex traffic situations and take various kinds of uncertainty into consideration, in a fashion that closely resembles human behavior and planning strategies. The advantages lie in the transparent choice of assumptions and, in particular, a transparent trade-off between conflicting goals, such as safety, comfort and pace. The risks that are tolerated to reach these goals can be chosen deliberately—defining the weights that trade off between safety and secondary goals is certainly difficult and expressly far beyond the scope of this paper, but it is assumed that once such a trade-off is technologically feasible, it should be made explicit.

Acknowledgments

This work was partially supported by Valeo Schalter und Sensoren GmbH within the V50 project, and by the Fraunhofer-Gesellschaft along with the state of Baden-Württemberg within their joint innovation cluster REM 2030.

References

[All07] E. Allen. *Modeling with Itô Stochastic Differential Equations*. Springer, 2007.

[BBF+08] Andrew Bacha, Cheryl Bauman, Ruel Faruque, Michael Fleming, Chris Terwelp, Charles Reinholtz, Dennis Hong, Al Wicks, Thomas Alberi, David Anderson, Stephen Cacciola, Patrick Currier, Aaron Dalton, Jesse Farmer, Jesse Hurdus, Shawn Kimmel, Peter King, Andrew Taylor, David Van Covern, and Mike Webster. Odin: Team VictorTango's entry in the DARPA Urban Challenge. *Journal of Field Robotics*, 25(8):467–492, 2008.

[Bul79] M. G. Bulmer. *Principles of Statistics*. Dover Publications, 2nd edition, 1979.

[GK02] Carl Geiger and Christian Kanzow. *Theorie und Numerik Restringiererter Optimierungsaugaben*. Springer-Lehrbuch Masterclass. Springer, 2002.

[GKP12] B. Goldengorin, V.A. Kalyagin, and P.M. Pardalos. *Models, Algorithms, and Technologies for Network Analysis: Proceedings of the First International Conference on Network Analysis*. Springer Proceedings in Mathematics & Statistics. Springer, 2012.

[Kum80] P. Kumaraswamy. A generalized probability density function for double-bounded random processes. *Journal of Hydrology*, 46(1-2):79–88, March 1980.

[KZP+08] Sören Kammel, Julius Ziegler, Benjamin Pitzer, Moritz Werling, Tobias Gindele, Daniel Jagzent, Joachim Schröder, Michael Thuy, Matthias Goebl, Felix von Hundelshausen, Oliver Pink, Christian Frese, and Christoph Stiller. Team annieWAY's autonomous system for the DARPA urban challenge 2007. In *International Journal of Field Robotics Research*, 2008.

[LTR12] Laura Leal-Taixé and Bodo Rosenhahn. *Pedestrian interaction in tracking: the social force model and global optimization methods*. Springer, September 2012.

[MBB+08] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhnke, D. Johnston, S. Klumpp, D. Langer, A. Levandowski, J. Levinson, J. Marcil, D. Orenstein, J. Paefgen, I. Penny, A. Petrovskaya, M. Pflueger, G. Stanek, D. Stavens, A. Vogt, and S. Thrun. Junior: The Stanford Entry in the Urban Challenge. *Journal of Field Robotics*, 2008.

[NW06] J. Nocedal and S. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, 2006.

[Øks03] B. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer, 2003.

[RZR+14a] M. Ruf, J. Ziehn, B. Rosenhahn, J. Beyerer, D. Willersinn, and H. Gotzig. Evaluation of an Analytic Model for Car Dynamics. In *International Conference on Mechatronics and Control (ICMC), Jinzhou, China*, pages 2446–2451, July 2014.

[RZR+14b] M. Ruf, J. Ziehn, B. Rosenhahn, J. Beyerer, D. Willersinn, and H. Gotzig. Situation Prediction And Reaction Control (SPARC). In Berthold Färber, Klaus Dietmayer, Klaus Bengler, Markus Maurer, Christoph Stiller, and Hermann Winner, editors, *9. Workshop Fahrerassistenzsysteme – FAS2014*, pages 55–66, March 2014.

[UAB+08] Christopher Urmson, Joshua Anhalt, Hong Bae, J. Andrew (Drew) Bagnell, Christopher R. Baker , Robert E Bittner, Thomas Brown, M. N. Clark, Michael Darms, Daniel Demitrish, John M Dolan, David Duggins, David Ferguson , Tugrul Galatali, Christopher M Geyer, Michele Gittleman, Sam Harbaugh, Martial Hebert, Thomas Howard, Sascha Kolski, Maxim Likhachev, Bakhtiar Litkouhi, Alonzo Kelly, Matthew McNaughton, Nick Miller, Jim Nickolaou, Kevin Peterson, Brian Pilnick, Ragunathan Rajkumar, Paul Rybski, Varsha Sadekar, Bryan Salesky, Young-Woo Seo, Sanjiv Singh, Jarrod M Snider, Joshua C Struble, Anthony (Tony) Stentz, Michael Taylor , William (Red) L. Whittaker, Ziv Wolkowicki, Wende Zhang, and Jason Ziglar. Autonomous driving in urban environments: Boss and the Urban Challenge. *Journal of Field Robotics Special Issue on the 2007 DARPA Urban Challenge, Part I*, 25(8):425–466, June 2008.

[Zie12] Jens R. Ziehn. Energy-based collision avoidance for autonomous vehicles. Master's thesis, Leibniz Universität Hannover, Germany, October 16th 2012.