

Title:

Handling Occlusions in Automated Driving Using a MEC  
Server-based Environment Model From Infrastructure Sensors

Authors:

Michael Buchholz, Johannes Müller, Martin Herrmann,  
Jan Strohbeck, Benjamin Völz, Matthias Maier, Jonas Paczia,  
Oliver Stein, Hubert Rehborn, and Rüdiger-Walter Henn

This article has been accepted for publication in IEEE Intelligent Transportation Systems Magazine. This is the author's version which has not been fully edited and content may change prior final publication.

Citation information: DOI 10.1109/MITS.2021.3089743

Copyright © 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

# Handling Occlusions in Automated Driving Using a MEC Server-based Environment Model From Infrastructure Sensors

Michael Buchholz\*, Johannes Müller\*, *Student Member, IEEE*, Martin Herrmann\*, *Student Member, IEEE*, Jan Strohbeck\*, *Student Member, IEEE*, Benjamin Völz†, Matthias Maier‡, Jonas Paczia§, Oliver Stein§, Hubert Rehborn§, and Rüdiger-Walter Henn‡

\*Ulm University, Institute of Measurement, Control and Microtechnology, Ulm, Germany

†Robert Bosch GmbH, Corporate Research, Renningen, Germany

‡Robert Bosch GmbH, Cross-Domain Computing Solutions, Leonberg, Germany

§Mercedes-Benz AG, Böblingen, Germany

**Abstract**—The on-board sensors' view of an automated vehicle (AV) can suffer from occlusions by other traffic participants, buildings, or vegetation, especially in urban areas. However, knowledge of possible other traffic participants in the occluded areas is crucial for an energy and comfort optimizing control of an AV. In such a case, information from infrastructure sensors sent via vehicle to anything (V2X) communication can help the AV. For such cases, we have developed and prototypically implemented a concept where an infrastructure environment model is generated from infrastructure sensors on a multi-access edge computing (MEC) server of an LTE/5G mobile network. This information extends the AVs' field of view and is beneficially integrated into their motion planning schemes. In this article, after a description of the modules of our approach, we present and discuss real-world results from our pilot site on a public junction with prototype AVs.

**Index Terms**—Connected Automated Driving, Field Test, Merging Scenarios, V2X

## I. BACKGROUND AND MOTIVATION

COMPARED to other highly automated systems, e.g. shop floor automation, where the environment is adapted or simplified to support the automation system, automated vehicles (AVs) in mixed traffic have to fulfill challenging tasks in very complex and changing environments, which are designed for human drivers. For example, the signalization of traffic rules by signs, traffic lights, or markings on the road targets the cognitive capabilities of human drivers. Thus, in AVs, video cameras and other sensors, like lidars and radars, collect the information around the vehicle as input for the algorithms responsible for driving [46]. In addition, different system

Part of this work was financially supported by the Federal Ministry of Economic Affairs and Energy of Germany within the program "Highly and Fully Automated Driving in Demanding Driving Situations" (project MEC-View, grant numbers 19A16010A, 19A16010B, and 19A16010I; project LUKAS, grant numbers 19A20004A, 19A20004C, and 19A20004F). Part of this work has been conducted within the ICT4CART project, which has received funding from the European Union's Horizon 2020 research & innovation program under grant agreement No. 768953. Content reflects only the authors' view and European Commission is not responsible for any use that may be made of the information it contains.  
 Corresponding author: M. Buchholz (phone: +49 731 50-27003; fax: +49 731 50-12-27003; e-mail: michael.buchholz@uni-ulm.de).



Fig. 1. Photo of the junction with the side road on the top right and the building occluding the priority road on the bottom right. The CAVs should turn right on the priority road like the white vehicle in the center. The building on the lower right occludes the CAVs' on-board sensors FOV.

approaches exist to communicate important information, e.g. the current state of a traffic light, to any connected vehicle [1]. A connected automated vehicle (CAV) can make use of, e.g., an ad-hoc network, like ITS-G5 (also called Dedicated Short Range Communication, DSRC), or a cellular mobile network, like LTE/5G. With that, it can retrieve important information from other road users (vehicle to vehicle, V2V) and/or other sources (vehicle to anything, V2X), like road-side units (RSUs). Besides communicating perception data, research has extensively studied approaches for cooperative trajectory coordination [41] as well, mostly by simulations. However, they are currently difficult to implement, as they typically assume all road traffic being CAVs [47].

Extending the previously mentioned approaches to communicate information from traffic infrastructure or other road users, in the German public funded project MEC-View<sup>1</sup>, we have conducted research on an approach that supports the

<sup>1</sup>Website: [www.mec-view.de](http://www.mec-view.de)

CAVs' on-board perception system with information from infrastructure sensors. Especially in urban areas, other traffic participants, houses, trees, etc., can heavily occlude the on-board sensors' field of view (FOV). Thus, we extend the CAV's FOV by communicating the information on traffic detected by infrastructure sensors. As automated driving poses demanding requirements on communication latencies, since up-to-date data are crucial for safe driving, our idea was to use a multi-access edge computing (MEC) server within a latency optimized mobile network. The server processes the sensors' data to an infrastructure environment model and predicts the movement of all objects for a few seconds to allow for predictive planning and to counteract latency. This can additionally reduce the computational effort on-board the CAVs.

To evaluate this approach in real traffic, a pilot installation has been built up at a suburban T-junction in the city of Ulm, Germany [3], see Fig. 1. There, a CAV approaching the junction on the side road has to give right of way when merging onto the priority road. However, the vehicles approaching the junction on the priority road are only visible to the on-board sensors when the CAV has already arrived at the yield line due to a building on the corner of the junction. Thus, the CAV has to slow down to (almost) stop before accelerating again even when there is a free main road.

Receiving the environment model including the objects' predictions from the MEC server, the CAV's motion planning can consider information from occluded areas as well as from areas beyond its own FOV. Thus, the vehicle can plan its own trajectory at the junction more smoothly, and the avoidance of braking and re-accelerating can be more energy efficient [27].

In this article, we concentrate on the end of the processing chain and compare two different methods of integrating the external environment model into the data processing of the CAV: One uses a track-to-track (T2T) fusion with the on-board environment model as a basis for a hierarchical motion-planning scheme. The other one solves the predictive motion planning holistically on the external information, while the reactive trajectory control uses the on-board environment model. Each approach has been implemented into one prototype CAV. We prove the feasibility of the overall system approach using results from test drives with the two prototype CAVs at our pilot site in real traffic. Based on that data, we evaluate the impact on the maneuver time of the CAVs as well as the influence on their energy consumption for different traffic scenarios and compare them to a baseline without external environment information.

## II. RELATED WORK

Traffic surveillance at junctions has already been reported in literature, as summarized in the survey [7]. However, many approaches concentrate on detecting and analyzing the traffic or, like [10], generating ground truth data for the evaluation of AV driving functions. Thus, the communication between infrastructure and CAVs is not in their focus and, if at all, often only regarded in a very generic way.

In many of those approaches that include a communication between infrastructure and the CAV, the sensor data or pro-

cessed information, like object detections, are communicated directly. In contrast, our work includes a mobile network featuring a MEC server, which provides fusion of the infrastructure sensors' object detections into an overall environment model as well as a prediction of these objects. This comes with the advantage that the computational load for this fusion and prediction tasks are not required in the vehicle [48]. Our approach can be extended to hybrid communication supporting cellular mobile networks as well as DSRC as shown in [4]. Other approaches, like [13], go further and move safety-relevant functionality at least partly into the infrastructure. In contrast, our approach assists the automated vehicle and can improve efficiency and comfort, while safety is ensured by the functionality on-board the vehicle, thus making it more resilient with respect to failures on communication or infrastructure side.

A first, but only very general overview on fusion and tracking algorithms for the junction surveillance application is already included in [7]. Today, Random Finite Sets (RFS) have become state-of-the-art for multi-object tracking. Based on the Finite State Set Statistics (FISST) theory [28], the Generalized Labeled Multi-Bernoulli filter has been the first closed form solution to the multi-object Bayes theorem that is computationally tractable [29]. Further extensions and approximations to reduce the computational effort and to allow for a real-time application have been developed, like the Labeled Multi-Bernoulli (LMB) filter [40]. Our approach presented in this paper uses a centralized LMB Multi-Object Tracker (MOT) [18], [19]. However, the application would also allow for a distributed implementation of the Bayes-optimal GLMB filter between the MEC server and the sensor processing units (SPUs), as we show in [17]. Alternatively, sub-optimal track-to-track schemes based on RFS are very popular, but suffer from unresolved challenges [5].

Concerning motion planning methods for the automated vehicles, the survey paper [47] includes an up-to-date summary. The methods can be clustered into graph-based, learning-based, and optimization-based methods. Graph-based methods discretize the state and action space and interpret the discretized states as nodes in a graph. Then, graph-search algorithms, such as the well-known A\* algorithm [9], [15] or rapidly exploring random trees [8], are used to determine the optimal path from the ego state to the AV's destination. Learning-based methods comprise deep reinforcement learning (deep RL) [16] and supervised end-to-end (E2E) learning approaches [23]. Deep RL tends to generalize better than E2E learning. However, both, graph-search methods and learning-based methods typically suffer from high computational burden. Optimization-based methods formulate the planning as an Optimal Control Problem (OCP) [14], [45], [49]. The OCP is then solved either analytically or numerically by model predictive control methods. Analytical solutions often solve the OCP up to some parameters, over which a grid search is conducted to find the overall optimal trajectory. Both approaches in this paper belong to this category.

The energy consumption of vehicles depends on traffic flow. An overview of different energy consumption models can be found in [2]. Some of the described models use current

movement information of vehicles (e.g., velocity and acceleration) and are especially interesting because these signals can be easily measured during field trials. Other approaches use static information from digital maps (e.g., positions of street signs or historical aggregated speed profiles) [22], [30]. However, these approaches are less flexible and more complex. Since we compare the energy consumption of drives with almost identical trajectories, in our case, the static setting is also identical and not significant to our relative comparison. To prove the traffic phase dependency of fuel consumption based on measured data from vehicles and Kerner's three phase traffic theory [24], Koller [25] introduced a consumption matrix which uses acceleration and velocity pairs. Rehborn *et al.* [38] use the energy matrix approach to compare the consumption of simulated automated and non-automated vehicles. The application of the energy matrix investigates the energy consumption with respect to traffic congestion using data from a microscopic simulation based on a stochastic Kerner-Klenov three-phase traffic flow model. The simulated data shows that traffic in a congested state requires up to 40% more energy (see also Rehborn *et al.* [39]). In contrast to these simulation results, this paper reveals the differences of energy consumption of CAVs with and without infrastructure support based on measured data from a field trial.

### III. SYSTEM OVERVIEW

Our main goal within the project MEC-View was to demonstrate proof of concept for an infrastructure sensor based assistance for CAVs with low latency. Additionally, to keep the computational effort on-board the vehicle as low as possible, a ready-to-use infrastructure environment model should be provided to the CAVs, including the objects' predictions for some seconds. Especially for the prediction task, this comes with the further advantage that the method can be highly adapted to the special characteristics and conditions of the local environment, while on-board the CAV, a generic method would have to be applied.

Although the main goal was a proof of concept, we have yet developed an system architecture that can serve as a basis for later commercial developments or further research. This architecture is shown in Fig. 2 and consists of the following components: The infrastructure sensors with SPUs, a mobile network (e.g. LTE or 5G), a MEC server with additional connection to backend services (e.g. for an up-to-date high definition (HD) map), and the CAVs. Except some control and error handling commands, which are neglected within the rest of this paper, the information flow is unidirectional from the sensors via the MEC server to the CAVs. Despite these simplifications for our proof of concept, our architecture can be easily extended for bilateral information flows, e.g. to incorporate information from the vehicles into the environment model, and to other means of communication, e.g. DSRC/ITS-G5, as shown in [4].

The sensors can be of different types, e.g. video cameras or lidar sensors, each being connected to an SPU. The SPUs perform an object detection on the sensors' raw data to reduce bandwidth requirements and transmission latencies. Based




Fig. 2. Overall system architecture.

on a calibration of the sensors with respect to some world coordinate frame, e.g. obtained from [34], and a synchronized time, e.g. from GNSS, the detected objects' information is communicated to the MEC server. The MEC server is also time-synchronized. As will be explained later in more detail, our generic object interface of the MEC server allows for independent and relatively cheap sensors of various types.

The MEC server expects the sensors to register via their SPU once they are ready to send data. Within that registration, the SPU transmits the frame rate for this sensor, i.e. the rate in which new sensor data are expected. Since an empty object message is sent if there are no detections, the MEC server can perform a watchdog functionality for each sensor. Additionally, the MEC server can discard messages received above some delay limit to avoid the use of outdated data. During sensor registration, the SPU also transmits the respective FOV. Thus, the MEC server can assess if a sensor extends the overall FOV, or if it delivers redundant information to some other sensors to improve accuracy. Furthermore, the MEC server uses an HD map to detect whether the area under supervision is insufficiently covered. If not, it switches to an error mode instead of sending unreliable information to the CAVs.

The CAVs also register at the MEC server if they want to receive the server's environment information. Then, the server supplies them with the environment model until they de-register.

In the following two sections, the main components of our system at the end of the processing chain, the infrastructure environment model and the motion planning in the CAVs using this model, are shortly summarized.

### IV. ENVIRONMENT MODEL FROM INFRASTRUCTURE SENSORS' DATA

The environment model comprises information about dynamic objects, e.g., vehicles, pedestrians, and cyclists, within the surveilled junction area. This information consists of a state estimation with dynamic information, like the position, orientation, and velocity of an object, as well as static information, like the class and extent of an object. Any state is modeled assuming a multi-variate normal error model, thus, a covariance matrix completes the state information. Further, an estimation of the existence probability and predictive information about the most probable future trajectories of the objects is appended. A detailed definition of the environment model can be found in [18].

#### A. Information Fusion and Multi-object Tracking

The environment model uses measurements of distributed sensors pre-processed by their associated SPUs, which are






Fig. 3. 2D rectangular shaped box model with the object reference points denoted by  $\oplus$  and schematic visualization of estimating length  $l$  and width  $w$  solely from two object reference point measurements marked by  $*$ .

fused on the MEC server using a centralized LMB MOT [18]. The interface corresponds to the environment model, which models dynamic objects using a rectangular shaped box approximation and with the above given static and dynamic state. However, the sensors are only required to measure one of the anchors of the rectangle, which are its center and its corners, as shown in Fig. 3, as well as a type and type probability. All other features are optional in the sensor interface. If an object is seen from multiple sensors with different perspectives, and thus different object reference points are reported, the developed LMB MOT is able to infer the object's extent as depicted in Fig. 3. This strategy allows a plug-and-play installation of a wide variety of different sensor kinds. As described in [18] in more detail, this is meant to address the price-sensitive mass-market for a potential later series application.

In cases where correct determination of an object reference point is not possible, e.g. if a lidar sensor cannot distinguish between front and rear of an object, an extended algorithm enables the object reference point association over time [19]. Hereby, multiple hypotheses are generated and thinned using statistical methods, which, however, increases computational cost of the method. An additional downstream model-based validation gate allows for simultaneous improvement of the computational effort as well as the performance of the filter by sorting out infeasible object states.

### B. Prediction

In order to support CAVs in their behavioral decisions further, a subsequent stage predicts the most probable future trajectories of all dynamic objects, using a straightforward map-based approach combined with simple process models, as depicted in Fig. 4. Hereby, dynamic objects are associated to appropriate lanes of a static HD map and predicted with a Kalman filter based on the most probable object class. A final stage associates a probability of occurrence based on the actual distance between track and appropriate lanes.

Although relatively simple, this approach is computationally very efficient and highly accurate if road participants adhere to the lanes. A yet more flexible approach based on state-of-the-art neural networks [43] could be used alternatively. For this purpose, a synthetic image of the environment model is compiled and fed to the network together with the current probabilistic state of the dynamic objects and their past




Fig. 4. Schematic visualization of the prediction method implemented on the MEC server: The measured object (red) is associated to the appropriate lane (blue) and predicted along it with an appropriately parametrized Kalman filter (shaded blue).

trajectories. This allows model-free inference of a set of unconstrained future trajectories, which are augmented with probabilistic measures. The method is proven to be highly accurate, since it won the international Argoverse Motion Forecasting Challenge 2019<sup>2</sup>.

## V. MOTION PLANNING FOR CAVs WITH EXTERNAL INFORMATION

One key concept of our overall system is leaving the responsibility of the decision making to the automated vehicles. Thus, the infrastructure sensors and environment model are not safety critical and can be realized much cheaper than if safety certification was necessary. In turn, we have developed a method [32] to assess the reliability of the incoming data on-board the CAV. This avoids critical driving situations due to erroneous data in the planning stage.

Whenever reliable external environment information is available, it can help the CAV to extend its FOV even beyond the normal limits. Accordingly, the motion planning of the CAVs has to be adapted to make use of this additional information. Within our work, we have developed two different methods, which we will explain after a short introduction into our reliability estimation scheme.

### A. Reliability Estimation

As shown later, reliable external data can increase the traffic efficiency in motion planning, whereas erroneous data may result in harm. Thus, the CAVs need to know whether the incoming data are reliable. The basic idea of estimating the reliability of the information is to use redundancy in the data to test for both, plausibility and consistency. If the data are plausible and consistent, this indicates high reliability, whereas low reliability is inferred otherwise.

Our reliability estimation scheme to govern this task bases on two consistency checks and two plausibility tests [32]:

- A prediction test checks if former predictions are consistent with current measurements;

<sup>2</sup>Website: <https://www.argoverse.org/tasks.html#forecasting-link>

- A map test compares whether or not the reported objects are consistent with the digital map of the CAV;
- An ego perception test judges the plausibility of the reported objects by comparing them with the objects from the ego perception;
- An ego localization test determines the data reliability by measuring how plausible the transmitted uncertainties (given in terms of covariance matrices) are, based on the reported position of the CAV.

An important challenge for the reliability estimation is that, usually, only very few measurements are available. Thus, the statistical uncertainty plays an important role. To cope with that, our method [32] bases on Subjective Logic (SL) [21], which is a theory to handle small sample sizes. With SL, we can describe the problem in terms of  $\beta$ -distributions representing the probability that the estimated reliability is correct, from which the p-value from classical statistics can be retrieved [31]. For possibly unreliable communication links, the reliability of the communication channel can also be assessed by our method [33]. This, however, can be omitted if the communication link is known to be stable.

### B. T2T Fusion-based Hierarchical Motion Planning

The first planning concept relies on a T2T fusion between the infrastructure environment model and the CAV's on-board environment model as input for a hierarchical motion-planning scheme. Although the CAV's sensors on their own are not able to observe the occluded priority road, they still provide the best possible state estimate for visible objects close to the CAV. As these objects are of highest importance for the CAV's short-time planning and safety evaluation, it can be mission critical to get the best possible data basis for them. Therefore, the T2T fusion essentially updates all objects seen by the infrastructure with the most recent state estimated by the CAV, minimizing any potential problems arising from aggregated infrastructure calculation and transmission latencies.

The main idea of the hierarchical planning scheme is to separate the problem into multiple sub-systems, which are easier to solve, interpret, and evaluate. Initially, we use the provided fusion output together with a highly precise map to identify priority traffic and relations between individual objects. This is possible due to the road geometry and priority rules provided by the map. Based on this information, the high-level planning component identifies gaps between priority objects and selects one of these gaps as merge target for the underlying motion planner. More details on this can be found in [44].

We predict the relevant objects as well as the gap between them to the intersection area by using a lane-based prediction similar to the one shown in Fig. 4. From this, we can implicitly define target variables for the following motion planner. Essentially, we compute a tuple comprising of target position, time, and velocity. Then, the motion planner provides a trajectory reaching this state. For this, we use a map-driven sampling-based approach to compute polynomial trajectories.

For simplification, we use the lane geometry provided by the map to define the lateral motion and then sample longitudinally




Fig. 5. Trajectory bundle for a merging scenario. The image shows the automated vehicle (black), objects (yellow), map (red), and a set of feasible trajectories (green). The trajectories use the z-axis to visualize the velocity. The selected trajectory is shown in blue.

in the 1D space along the lane. Usually, the 1D sampling is rather straightforward, as we only have to match a target velocity and possibly a position, e.g. when approaching a stop point. Typically, the time when matching has to be achieved is not defined. However, for the merging scenario, we do not have this degree of freedom anymore, since the merging time is relevant to meet the gap. Therefore, we sample piecewise polynomials to reach the given end condition within our vehicles kinodynamically constrained motion space. We use a parametrized sampling range to sample different intermediate velocities around the provided target velocity, while considering the CAV's current velocity. Essentially, our piecewise trajectories follow the scheme: accelerate/decelerate towards the intermediate velocity, hold it for some time and finally decelerate/accelerate towards the target velocity. To limit the required computational power, we heavily rely on fast and lightweight pre-estimations to calculate a feasible set of intermediate sampling points.

After the 1D sampling, we transform the trajectories into the 2D space. After a violation check with respect to our parametrized kinodynamic constraints, we rate all sampled trajectories by a pre-defined cost function. As this is done in the 2D space, we can account for lateral safety (constraints) and comfort (costs) measures, guaranteeing, e.g., that we will not drive with maximum velocity into a sharp curve. Altogether, the cost function will select a trajectory that minimizes the required time by maximizing the velocity, while ensuring to minimize both longitudinal and lateral acceleration and jerk. Figure 5 shows a sampled trajectory bundle for our merging scenario.

### C. Fusion-free Motion Planning

Particularly in situations with heavy occlusion, the T2T fusion of the external and ego environment model often reduces to a concatenation of object lists and, thus, loses its merits. Additionally, the ego perception is designed to feature sufficient accuracy for the motion planning. Hence, the motion planning might not profit from a possible increase in accuracy by T2T fusion.

The key idea of the fusion-free planning is to use the external and the ego environment model in parallel for decision making [35]. Figuratively speaking, the planning uses the infrastructure data to predictively synchronize the CAV's

motion such that it can merge into an oncoming traffic gap, while the ego perception is used for reactive planning in case of other traffic participants occurring on the road ahead.

As a first step, the planner samples suitable target points according to the map and the dynamic situation context derived from the environment models. In contrast to the hierarchical planner, this planner samples in 1D over time rather than position. The retrieved target points also comprise target positions, velocities, and time. Due to the sampling approach, the planning can easily consider multiple behavior options in parallel and solve motion planning and behavior planning holistically.

We formulate the trajectory planning as an OCP. The OCP's cost function accounts for jerk and the end time of the maneuver. One specific effect that occurs when dealing with external environment models containing uncertainties is that the planning decision might frequently change while approaching the intersection. To mitigate this effect, we consider a time weighting of the jerk in the OCP's cost function. With that, the planning favors smoother trajectories in the near future at the expense of potentially higher jerks far into the future. Our approach allows calculating analytical solutions of the OCP for the CAV's motion from its ego state to the target states as trajectory candidates.

Furthermore, the planning comprises a risk model, which describes the residual risk of a planned trajectory and its corresponding behavior. In a next step, we calculate the residual risk for each trajectory candidate and add it to its dynamic cost. All candidates exceeding the maximum accepted residual risk, which can be chosen arbitrarily, are discarded. For the extreme choice of allowing zero residual risk, the safety of the motion planning can be guaranteed formally. The remaining candidates are checked for possible violations of kinodynamical constraints. From this final set of candidates, the one with minimal cost is chosen. A fail-safe strategy ensures that there is always at least one feasible trajectory. A more detailed description of the fusion-free motion planning approach can be found in [35].

## VI. EXPERIMENTAL SETUP

One important goal of our work was to prove and evaluate our approach in real traffic. To do so, we have implemented the overall functionality of the system architecture on a pilot infrastructure installation and into two prototype CAVs, which are described in more detail in the following. Additionally, we shortly describe some regulatory needs and organisational measures, e.g., to account for the prototype characteristic of our setup during the experiments.

### A. Pilot Site for Evaluation in Real Traffic

Together with further partners within the project MEC-View, we built up a pilot site at an unsignalized T-junction in a suburban area of the city of Ulm, Germany [3]. As shown in Fig. 6, a side road ends at the priority road, and vehicles entering the priority road have to give right-of-way. In this work, we concentrated on merging scenarios for CAVs when turning right onto the priority road, as shown in Fig. 1. As in

many urban areas, a building on the left at the end of the side road occludes the on-board sensors' view on the oncoming traffic on the priority road nearly up to the yield line. Thus, our pilot site comprises of video cameras and lidar sensors surveilling the oncoming roads close to the T-junction. The approximate FOVs on the side road and the priority road are marked by blue and yellow areas in Fig. 6, respectively. Most of the sensors have been mounted to five existing lampposts along the roads. For some of the lidar sensors and video cameras, two temporary poles have been placed on concrete bases in the junction area, located on the opposite sides of the street light poles. As can be seen in Fig. 1, they are interconnected with one small auxiliary pole at the corner of the occluding building via traverses, which are used for cable routing and stability. One SPU per lamppost as well as one SPU for the temporary poles have been placed in outdoor housings on the sidewalks to avoid high wind loads except for one case, where mounting on the pole was possible. They serve for data acquisition, object detection and communication to the MEC server via the mobile network with a frequency of 10 Hz for the cameras and 20 Hz for the lidars. For the SPUs at our pilot installation, we use standard PCs with a graphic processing unit (GPU) for object detection on camera images and a USB modem with external antenna to connect to the mobile network. The sensors are connected to their respective SPU via cables, where multiple sensors can be connected to one SPU.

For this work, we use the object detections from five monocular cameras and two lidars with 16 static beams mounted on the lampposts only. This simple setup was chosen on purpose to be close to a possible series setup in the future. Both sensor types are restricted in their sensory abilities. The lidar sensors have a highly restricted field of view, such that the lidar beams never capture large objects completely, but only their longitudinal profile. The monocular cameras, in contrast, suffer from an ambiguity when transforming from the image coordinate system to a real world Cartesian coordinate system, so that an object's center cannot be measured perfectly. These problems are tackled by the ability of our generic interface to allow position-only detections of one of the pre-defined reference points of an object. The sensors used for this work yield an average position error in the range of 0.34 m to 1.54 m. Those errors comprise a highly position-dependent systematic error that has not been corrected. For the cameras, this error mainly stems from the homography that transforms the two-dimensional camera coordinates to the three-dimensional Cartesian coordinates and assumes a flat ground, which does not perfectly hold for the junction area. For the lidar sensors, the error increases with higher distances due to the beam widening.

A latency-optimized test mobile network then connects the sensors' SPUs, which perform the object detections, with the MEC server. The same network connects the MEC server with the prototype CAVs. For the ease of implementation for the proof of concept, we have used a proprietary message format between the MEC server and the CAVs. However, ETSI messages like the "Collective Perception Message (CPM)" for the environment model data are in the process of standardization




Fig. 6. Drawing of the pilot site with approximate FOV of the infrastructures on the priority road (yellow) and on the side road (blue), CAV (red), and on-board FOV (green).

and could also be used, as described in [4].

### B. Connected Automated Vehicles

For the experiments at our pilot junction, we have prepared two prototype CAVs, one (CAV#1) operated by Robert Bosch GmbH, the other one (CAV#2) by the Institute of Measurement, Control and Microtechnology at Ulm University. To cover a broad spectrum of possible parameters and implementations, the CAVs strongly differ in their setup and parametrization.

1) *CAV#1 operated by Robert Bosch GmbH*: In this project, we use an existing test vehicle of Bosch Research. The vehicle is equipped with a Velodyne HDL-64 lidar, five Bosch radars, and a Bosch stereo video camera for object detection, as well as a GeneSys ADMA for highly precise ego localization based on a real-time kinematic (RTK) system with differential GPS (D-GPS) and gyro sensors. Additionally, we mounted a Bosch Car Communication Unit (CCU) to enable communication with the MEC-Server. For the environment modeling, we fuse all on-board sensors into a vehicle-centric environment model. Afterwards, track-to-track fusion is used to incorporate the MEC information [11]. For our experiments, the T2T fusion-based hierarchical motion planning is applied in CAV#1. A sketched processing architecture and a picture of the vehicle are presented in Fig. 7. The vehicle features a very strong motorization. Accordingly, the parametrization of both, the planning and control have been chosen to cover a sporty driving style. Hence, the overall planning objective is the travel time minimization by allowing higher vehicle dynamics where possible. In the trade-off between time and energy efficiency, fuel consumption is weighted less.

2) *CAV#2 operated by Ulm University*: The actual CAV of Ulm University settles on the basics described in [26]. However, the CAV itself is a different vehicle and many subsystems




Fig. 7. Image and architecture of the CAV#1 operated by Bosch Research.

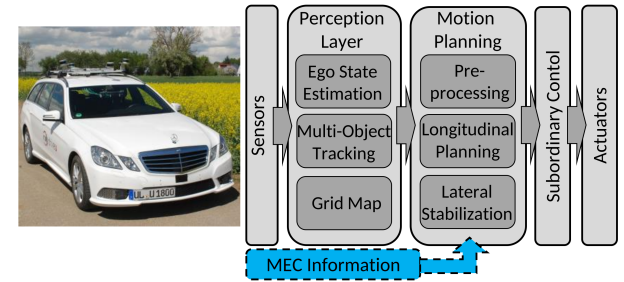


Fig. 8. Image and architecture of the CAV#2 operated by Ulm University.

experienced major changes and updates. The sensor setup was updated and, in this project, we use a Velodyne VLP-32 lidar and three Continental ARS 408 radars. Furthermore, the environment model [12] uses an efficient implementation of the dynamic occupancy grid map (DOGMA) [36]. The pre-processing algorithms for the sensors use state-of-the-art detectors based on modern machine learning methods [6], [20]. Moreover, the complete software stack was rewritten and migrated to the ROS framework [37]. Finally, the AV has been extended by the same Bosch CCU used for CAV#1 to allow for communications with the MEC server. CAV#2 uses the fusion-free motion planning in our experiments. An image of this vehicle as well as an architecture overview are given in Fig. 8. With respect to the parameterization, a more conservative behavior is chosen, leading to an increased fuel efficiency at the cost of increased travelling time.

### C. Organisational and Regulatory Aspects

The two CAVs used in this project both have an approval as automated test vehicles for public roads in Germany, which requires an approval of the mechanical changes of the vehicle, if, e.g., sensors have been integrated in potential crash zones, as well as an approved safety concept including technical measures, like the take-over systems described below, and organisational measures, like training for the drivers and obligatory technical tests before driving. As a result, this approval comes with several requirements for the automated operation. Since the systems and software for automation did not undergo a homologation process, a trained safety driver is required to supervise the driving behavior and has to be ready to take over the control any time. This take-over can be initialized either by moving the steering wheel manually or by pressing any of the pedals. Additionally, the automation mode can be ended by an emergency switch. Before starting test drives, the safety driver



has to perform a take-over functionality test during standstill. Furthermore, a co-driver is required to assist the driver, e.g., by monitoring the automation system. Both, the safety driver and the co-driver must be from the staff of the owning company or institution and have to be familiar with the specific vehicle. Thus, although testing L4 [42] automated driving functions on public roads, from a legal perspective, the CAVs are still seen as human driven with the safety drivers being responsible for the vehicles' maneuvers.

The installation and operation of a pilot site in real traffic includes various organisational and regulatory aspects, from which three are discussed in more detail. First, a permission of the infrastructure owner, in our case, from the city of Ulm, is required to operate the pilot site. The city has been involved in our planning already since the proposal phase of the project and supported our work as an associated partner of the MEC-View project. Thus, the following requirements to retrieve such a permission for our installation could directly be considered in the design phase. All equipment installed on the sidewalks had to be placed with a minimum distance of 0.5 m to the curb and marked with reflecting warning signs. A clear width of at least 1 m had to remain on the sidewalk after placing SPU housings and the temporary poles' installation, which also should not interfere with usual pedestrian routes. The traverses of the temporary poles leave a clear height of at least 6.5 m to not restrict the passage even for high vehicles. Finally, the power supply is realized via the street light supply to avoid underground constructions. The city of Ulm has changed the electrical connection of the street lights such that one phase is freely available and always on for our equipment at the light poles, while some of the poles share the same phase. Thus, the components, especially the SPUs, had to be chosen carefully to not exceed the maximum allowed overall power.

Secondly, data protection must be taken care of, particularly, since video images of traffic participants are processed by the infrastructure sensors and the CAVs. For the protection of personal data, each of the partners involved in their processing either already had or developed a data protection concept according to the respective laws. Additionally, a common explanation on our data processing was published on the project's website and all residents close to the pilot site have been informed by a letter and by publications in local media. Finally, pictograms as well as information displays with a link to the project's website and contact information also make other traffic participants aware of the data processing and further information sources.

Thirdly, special safety measures have been taken to account for the prototype status of the implemented system. To reach a reasonably high safety level, each component of the overall system has been tested in simulation and preliminary tests before employing it to the real world experiment. Still, the implementation is only a prototype, lacking sophisticated automated diagnosis and failure detection measures. For example, there is no tracking of possible sensor degradation or dynamic adaptation of the sensors' FOV. Also, we have not implemented security measures like certificates to ensure the identity of communication partners, since the tests are performed in a private test mobile network. For our proof-of-concept,

no standardized testing like for the safety of the intended functionality (SOTIF) has been addressed so far. Therefore, for our test drives, we have voluntarily developed a safety concept to reduce the risk of possible harm to any person in this public area. Like for the CAVs, this safety concept relies on trained humans to supervise the prototypical infrastructure system. Since the CAVs should merge into a gap not visible also to the safety drivers most of the time when approaching the T-junction, we additionally included a safety marshal at the junction. This safety marshal assesses the situation on its own and signalizes clearance to the safety drivers. Finally, we decided to have informed persons as drivers in the vehicles at the back end of the gap. This is especially relevant when considering the fact that human drivers would not expect any vehicle entering the junction from the side road at high speed due to the occlusions at the junction.

## VII. EVALUATION

This section discusses the experimental results of our approach for infrastructure support of CAVs at our pilot site. While the first two subsections shortly present evaluations of the environment model and the reliability estimation, respectively, the further sections contain the evaluation of the overall system performance. For the latter, we performed several test drives with both available CAVs to collect data in real traffic. The test scenarios as well as the performance measures have been designed and selected by ourselves based on the intended functionality of our system and the research objectives of our project, since there have been no regulations on specific tests.

### A. Environment Model Evaluation

Independent from the test drives for the overall system assessment, we have first evaluated the environment model. This model is fed with the object detections of the five monocular cameras and the two lidars via our own highly generic interface. We have evaluated the environment model on data from manual test drives with the CAVs using their precise localization as ground truth. In that data, our environment model shows an average position error of 1.14 m. It especially suffers from the biased measurements of the sensors at positions where complimentary sensors show inconsistent biases. However, since the infrastructure's data are used in addition to the CAVs on-board perception, the uncertainty in the infrastructure environment model can be mitigated. In the motion planning, the uncertainties are accounted for in terms of safety margins.

### B. Reliability Estimation Evaluation

The second evaluation performed independent from the overall test drives addresses the new reliability estimation method, which is performed on data from a subset of the sensors at the pilot site. The method has been assessed on 14 sequences, which have been manually classified reliable or unreliable. Figure 9 presents the respective results. The reliability on the  $x$ -axis of Fig. 9 refers to the first order probability that the received external information actually is




Fig. 9. Evaluation of the reliability estimation scheme. It shows that for all sequences where the infrastructure worked as expected, the estimated reliability has a  $p$ -value above 90% for a confidence level of 90%.

correct. In turn, the statistical confidence of the  $y$ -axis is a second order probability reflecting the statistical evidence, which supports the assumed reliability. The evaluation shows with almost certain confidence that no unreliable sequence is rated with an estimated reliability above 70%. In turn, all reliable sequences reach at least 90% estimated reliability with a confidence above 90%. Hence, the two classes *reliable* and *unreliable* are easily separable. Furthermore, it has been evaluated in [31] that the reliability of the reliability estimation scheme reaches a  $p$ -value of 95%. Accordingly, the proposed method is well suited to further enhance the SOTIF.

### C. Overall System Latency

Now, the results from the overall system evaluations are reported. First, we analyzed the overall latency from the sensor measurements until the CAVs receive the infrastructure environment model. This latency includes the pre-processing of the data on the SPUs, the delay of the communication to the MEC server, the processing time on the MEC server, and, finally, the communication delay between the MEC server and the vehicle. The analysis of more than 4300 environment model messages received during the test drives results in an average latency of 242 ms (minimum 152 ms, maximum 627 ms). The number is higher than typical latencies on-board CAVs, which we estimate to about 150 ms for similar processing steps.

However, this has several reasons besides the accumulated communication delays of approximately 20 ms to 30 ms. Due to the high mounting positions, the video cameras potentially see a much higher number of objects. Additionally, the overall FOV at the pilot site is much larger than that of a CAV, leading to a higher number of objects to keep track of; and the generic interface with object reference points requires more computational effort, e.g., to resolve ambiguities. Finally, the buffering mechanism of the fusion algorithm causes additional variance in the overall latency, since the sensors' capturing are not necessarily synchronized. Considering all those, the overall system performs very well with respect to latency, which is additionally mitigated by the provided objects' predictions.

### D. Traffic Efficiency

To assess a possible gain in traffic efficiency for the CAVs, we have clustered the test drives into different maneuver

TABLE I  
EVALUATION DATA FROM TEST DRIVES WITH THE TWO CAVS AT THE PILOT SITE IN REAL TRAFFIC. THE NUMBERS BETWEEN THE TWO CAVS VARY IN GENERAL, SINCE THESE ARE OF DIFFERENT TYPE. A SIGNIFICANT EFFICIENCY GAIN RESULTS FOR BOTH CAVS.

Maneuver	CAV#1	CAV#2
<b>Free lane on priority road (baseline w/o infrastructure)</b>		
Maneuver time	12.6 s	17.7 s
<b>Free lane on priority road</b>		
Average no. of objects in V2X env. model	3.6	4.1
Maneuver time	8.3 s	10.9 s
<b>Merging behind another vehicle</b>		
Number of evaluated maneuvers	6	3
Average no. of objects in V2X env. model	3.5	4.9
Maneuver time (mean)	11.0 s	13.6 s
Maneuver time (minimum)	9.4 s	11.7 s
Maneuver time (maximum)	12.4 s	15.7 s
<b>Merging into a gap</b>		
Number of evaluated maneuvers	5	8
Average no. of objects in V2X env. model	4.0	5.2
Maneuver time (mean)	10.7 s	12.2 s
Maneuver time (minimum)	9.1 s	10.8 s
Maneuver time (maximum)	13.1 s	15.3 s

classes. Table I summarizes the results. To give an idea of the complexity of the situations, we noted the average number of relevant objects in the infrastructure environment model during the maneuvers. As a measure for efficiency, we use the maneuver time measured from 40 m before to 20 m after the yield line. The numbers between the two CAVs vary in general, since these are of different type with different vehicle dynamics and a different parametrization as described in Section VI-B. Because the pilot site comprises public roads with real traffic, test conditions vary and are not identically reproducible for multiple tests. Thus, merging when turning right onto the priority road with no other traffic on the target lane is the only class with an approximately deterministic behavior, where a fair comparison with a baseline experiment without external information is possible. As can be seen in Table I, both CAVs significantly gain from the knowledge of a free lane, since they do not need to (almost) stop at the yield line like in the case without infrastructure sensing. As expected from the more sporty parametrization, CAV#1 is more dynamic and, thus, has a smaller total travelling time than CAV#2. The more conservative parametrization of CAV#2 leads to smaller accelerations and decelerations and, thus, longer traveling times. As a result, this vehicle can gain even more from the information provided by the infrastructure.

The remaining evaluated test drives constitute the two maneuver clusters "merging behind a vehicle" and "merging into a gap". The varying times within one category result mainly from the variation of the actual traffic situations on the priority road. Note that without infrastructure sensors, the vehicle would also have to drive almost to the yield line before being able to see the oncoming traffic situation on the priority road in both cases. To avoid any critical situation, the CAVs would slow down as in the baseline scenario. Thus, with the environment model from the infrastructure, the CAVs make use of the extended FOV also in these cases and plan their




Fig. 10. Visualization of the empirical fuel consumption matrix with the parameters velocity and acceleration from [25]. Color indicates normalized consumption per second.

motion in coordination with the oncoming traffic, as can be seen from the maneuver times. This improves the efficiency of the traffic flow and, additionally due to less braking and acceleration, the comfort on-board the CAVs.

#### E. Energy Consumption

The energy consumption is computed using an empirically measured and normalized consumption matrix [25], which was built from 90 h driving data collected in a mid-size vehicle with a combustion engine by different drivers in various traffic situations. In the consumption matrix, for each velocity-acceleration pair, a normalized consumption value is assigned. The consumption is stated as value per second and normalized in the interval  $[0, 1]$ , i.e., it has no unit and the maximum consumption value is 1. These normalized consumption values are sufficient to determine relative consumption differences between two drives along comparable trajectories. The energy consumption of each drive is computed by accumulating the normalized consumption values of all measured acceleration-velocity pairs. We choose  $2 \text{ km/h}$  as interval width for the velocity and  $0.1 \text{ m/s}^2$  for the acceleration. Figure 10 shows a heat map visualization of the normalized fuel consumption matrix. Consumption values near zero are visualized in dark blue, while red elements indicate consumption values close to the maximum. Acceleration-velocity pairs without a measured consumption value are depicted in white. It can be seen that the energy consumption is relatively small during deceleration and increases strongly at  $50 \text{ km/h}$  for accelerations between  $1 \text{ m/s}^2$  and  $3.5 \text{ m/s}^2$ .

In Table II, the relative energy consumption values of all test drive categories are shown with respect to the baseline. Like for the traffic efficiency evaluation, we use the test drive without infrastructure support as baseline. The different parametrization and motorization of the two CAVs (cf. Section VI-A) lead to different behaviors regarding velocity and acceleration, as shown in Fig. 11 and Fig. 12, respectively. Comparing the average velocities of each category, it is notice-

TABLE II  
ENERGY CONSUMPTION IN EACH CATEGORY RELATIVE TO THE CORRESPONDING BASELINE MANEUVERS OF EACH CAV. THE DIFFERENCES BETWEEN CAV#1 AND CAV#2 STEM FROM DIFFERENT PARAMETRIZATIONS.

Maneuver	CAV#1	CAV#2
<b>Free lane on priority road (baseline w/o infrastructure)</b>		
Consumption	100%	100%
<b>Free lane on priority road</b>		
Consumption (mean)	121%	59%
<b>Merging behind another vehicle</b>		
Consumption (mean)	119%	85%
Consumption (minimum)	111%	55%
Consumption (maximum)	135%	104%
<b>Merging into a gap</b>		
Consumption (mean)	121%	69%
Consumption (minimum)	72%	54%
Consumption (maximum)	145%	87%




Fig. 11. Exemplary velocity and acceleration profiles of the CAV#1.

able that the CAV#1 drives between  $1.4 \text{ km/h}$  and  $5.2 \text{ km/h}$  faster than CAV#2. This difference is especially noticeable in the baseline (blue), where CAV#1 only decelerates to  $10 \text{ km/h}$  and avoiding a full stop, while CAV#2 comes to a standstill.

The results show that CAV#1 consumes about 20% more energy on average in all infrastructure supported categories compared to the baseline. This is due to its more sportive parametrization (cf. Section VI-A) for merging, which decreases the travel time as shown in Section VII-D. This effect is further intensified by the fact that the reference run contains a particularly comfortable and thus energy-efficient braking towards the stop line, as can be seen from the blue lines in Fig. 11. In contrast, with CAV#2, almost all infrastructure supported maneuvers require less energy than the baseline due to the higher priority of energy efficiency in its parametrization. In many cases, this saving is significantly high and comes with additional lower traveling times as reported in Section VII-D. For both CAVs, it can be seen from the variations in the velocity and acceleration profiles that the complexity of the categories “merging behind another vehicle” and “merging into a gap” is much more complex, because time and velocity have to be adapted to reach the priority




Fig. 12. Exemplary velocity and acceleration profiles of the CAV#2.

road at the right time with the required target velocity. For the example of merging behind another vehicle given in Fig. 11 for CAV#1, the vehicle accelerates first and then decelerates when approaching the junction. This can be explained by a change in the velocity of the vehicle on the priority road. At first, merging before the other vehicle seemed possible due to its sporty parametrization, but after the velocity of the other vehicle changed, the CAV#1 has to adapt its behavior in order to prevent a collision and merges after the other vehicle. The results of this work confirm the simulation results from [38], according to which an infrastructure supported AV is able to save about 25% of travel time and up to 50% of the energy consumption.

## VIII. CONCLUSION

In this paper, we proposed a communication-based approach using an external environment model from infrastructure sensors to extend the FOV for motion planning of CAVs at occluded intersections. We described the main components of our proposed architecture, which have been implemented on a pilot site in public traffic. Two motion planning algorithms using the external information have been developed and implemented in two prototypical CAVs. Test drives at the pilot junction have demonstrated proof of concept of our approach. The evaluation of these overall system tests showed very good latency results, and both motion planning algorithms could benefit greatly from the external information, improving traffic efficiency and passenger comfort when merging onto the priority road at the occluded T-junction. For the defensive parametrized CAV#2, an additional reduction of the energy consumption could be shown, while the sporty parametrization of CAV#1, despite its considerably shorter time in the baseline, could improve traveling times even further. Future work will address the quality of the environment model and the underlying sensor data pre-processing. Additionally, we will conduct research on cooperative behavior of the traffic participants locally managed on the MEC server.

## REFERENCES

[1] K. Abboud, H. A. Omar, and W. Zhuang, "Interworking of DSRC and Cellular Network Technologies for V2X Communications: A Survey,"

*IEEE Transactions on Vehicular Technology*, vol. 65, no. 12, pp. 9457–9470, 2016.

[2] R. Akçelik, C. Bayley, D. Bowyer, and D. Biggs, "A hierarchy of vehicle fuel consumption models," *Traffic Engineering and Control*, vol. 24, pp. 491–495, 10 1983.

[3] M. Buchholz, M. Herrmann, J. Müller, V. Belagiannis, P. Pavlov, B. Hättig, S. Schulz, and R.-W. Henn, "A Digital Mirror: A Mobile Edge Computing Service Based on Infrastructure Sensors," in *ITS World Congress Copenhagen*, 2018.

[4] M. Buchholz, J. Strohbeck, A.-M. Adaktylos, F. Vogl, G. Allmer, S. Cabrero Barros, Y. Lassoued, M. Wimmer, B. Hättig, G. Massot, C. Ponchel, M. Bretin, V. Sourlas, and A. Amditis, "Enabling automated driving by ICT infrastructure: a reference architecture," in *8th Transp. Res. Arena TRA 2020 (conf. cancelled)*, 2020. [Online]. Available: <https://oparu.uni-ulm.de/xmlui/handle/123456789/26086>

[5] A. Buonviri, M. York, K. LeGrand, and J. Meub, "Survey of Challenges in Labeled Random Finite Set Distributed Multi-Sensor Multi-Object Tracking," in *Proceedings of the 2019 IEEE Aerospace Conference*, vol. 2019-March. IEEE, 3 2019, pp. 1–12. [Online]. Available: <https://ieeexplore.ieee.org/document/8742216/>

[6] A. Danzer, T. Griebel, M. Bach, and K. Dietmayer, "2D Car Detection in Radar Data with PointNets," in *Proceedings of the 2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019.

[7] S. R. E. Datondji, Y. Dupuis, P. Subirats, and P. Vasseur, "A Survey of Vision-Based Traffic Monitoring of Road Intersections," *IEEE Trans. Inf. Theory*, vol. 17, no. 10, pp. 2681–2698, 2016.

[8] M. Du, T. Mei, H. Liang, J. Chen, R. Huang, and P. Zhao, "Drivers Visual Behavior-guided RRT Motion Planner for Autonomous On-road Driving," *Sensors*, vol. 16, no. 1, p. 102, 2016.

[9] D. Fassbender, B. Heinrich, and H. Wuensche, "Motion Planning for Autonomous Vehicles in Highly Constrained Urban Environments," in *Proc. IEEE/RSJ Int. Conf. on Intell. Robots Systems (IROS)*, Daejeon, South Korea, Oct. 2016, pp. 4708–4713.

[10] T. Fleck, K. Daaboul, M. Weber, P. Schörner, M. Wehmer, J. Doll, S. Orf, N. Sußmann, C. Hubschneider, M. R. Zofka, F. Kuhnt, R. Kohlhaas, I. Baumgart, R. Zöllner, and J. M. Zöllner, "Towards Large Scale Urban Traffic Reference Data: Smart Infrastructure in the Test Area Autonomous Driving Baden-Württemberg," in *Intelligent Autonomous Systems 15*, M. Strand, R. Dillmann, E. Menegatti, and S. Ghidoni, Eds. Cham: Springer International Publishing, 2019, pp. 964–982.

[11] M. Gabb, H. Digel, T. Müller, and R.-W. Henn, "Infrastructure-supported Perception and Track-level Fusion using Edge Computing," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1739–1745.

[12] F. Gies, J. Posselt, M. Buchholz, and K. Dietmayer, "Extended Existence Probability Using Digital Maps for Object Verification," in *Proceedings of the 2020 IEEE 23rd International Conference on Information Fusion (FUSION)*. IEEE, 2020.

[13] S. Gopalswamy and S. Rathinam, "Infrastructure Enabled Autonomy: A Distributed Intelligence Architecture for Autonomous Vehicles," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 986–992.

[14] F. Gritschneider, K. Graichen, and K. Dietmayer, "Fast Trajectory Planning for Automated Vehicles Using Gradient-Based Nonlinear Model Predictive Control," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Madrid, Spain, Oct. 2018, pp. 7369–7374.

[15] P. Hart, N. Nilsson, and B. Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," *IEEE Trans. Sys. Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[16] A. Haydari and Y. Yilmaz, "Deep Reinforcement Learning for Intelligent Transportation Systems: A Survey," *IEEE Trans. Intell. Transp. Syst.*, pp. 1–22, 2020.

[17] M. Herrmann, C. Hermann, and M. Buchholz, "A Distributed Generalized Labeled Multi-Bernoulli Filter for Synchronous Multi-Sensor Systems," *IEEE Trans. Signal Process.*, 2021, under review.

[18] M. Herrmann, J. Müller, J. Strohbeck, and M. Buchholz, "Environment Modeling Based on Generic Infrastructure Sensor Interfaces Using a Centralized Labeled-Multi-Bernoulli Filter," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019, pp. 2414–2420.

[19] M. Herrmann, A. Piroli, J. Strohbeck, J. Müller, and M. Buchholz, "LMB Filter Based Tracking Allowing for Multiple Hypotheses in Object Reference Point Association," in *2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, 2020, pp. 197–203.

[20] M. Herzog and K. Dietmayer, "Training a Fast Object Detector for LiDAR Range Images Using Labeled Data from Sensors with Higher Resolution," in *Proceedings of the 2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 2020.

[21] A. Jøsang, *Subjective Logic*. Springer International Publishing, 2016.



- [22] D. Karbowski, S. Pagerit, and A. Calkins, "Energy Consumption Prediction of a Vehicle along a User-Specified Real-World Trip," *26th Electric Vehicle Symposium 2012, EVS 2012*, vol. 3, 12 2012.
- [23] A. Kendall, J. Hawke, D. Janz, P. Mazur, D. Reda, J.-M. Allen, V.-D. Lam, A. Bewley, and A. Shah, "Learning to Drive in a Day," in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2019, pp. 8248–8254.
- [24] B. S. Kerner, *The Physics of Traffic*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004.
- [25] M. Koller, "Kraftstoffmehrverbrauch auf Basis dynamischer Verkehrslagen," Ph.D. dissertation, Universität Tübingen, 2015.
- [26] F. Kunz, D. Nuss, J. Wiest, H. Deusch, S. Reuter, F. Gritschneider, A. Scheel, M. Stübler, M. Bach, P. Hatzelmann, C. Wild, and K. Dietmayer, "Autonomous driving at Ulm University: A modular, robust, and sensor-independent fusion approach," in *2015 IEEE Intelligent Vehicles Symposium (IV)*, 2015, pp. 666–673.
- [27] M. Lederer, *Energieeffizientes Fahren*. Wiesbaden: Springer Fachmedien Wiesbaden, 2014, pp. 307–321. [Online]. Available: [https://doi.org/10.1007/978-3-658-04451-0\\_8](https://doi.org/10.1007/978-3-658-04451-0_8)
- [28] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood: Artech House Inc., Norwood, 2007.
- [29] R. P. S. Mahler, *Advances in Statistical Multisource-Multitarget Information Fusion*. Norwood: Artech House Publishers, 2014. [Online]. Available: [https://www.ebook.de/de/product/22325050/ronald\\_p\\_s\\_mahler\\_advances\\_in\\_statistical\\_multisource\\_multitarget\\_information\\_fusion.html](https://www.ebook.de/de/product/22325050/ronald_p_s_mahler_advances_in_statistical_multisource_multitarget_information_fusion.html)
- [30] C. F. Minett, A. M. Salomons, W. Daamen, B. van Arem, and S. Kuijpers, "Eco-routing: Comparing the fuel consumption of different routes between an origin and destination using field test speed profiles and synthetic speed profiles," in *2011 IEEE Forum on Integrated and Sustainable Transportation Systems*, 2011, pp. 32–39.
- [31] J. Müller and M. Buchholz, "Subjective Logic Reasoning: an Urn Model Intuition and Application to Connected Automated Driving," at - *Automatisierungstechnik*, vol. 69, no. 2, pp. 111–121, 2021.
- [32] J. Müller, M. Gabb, and M. Buchholz, "A Subjective-Logic-based Reliability Estimation Mechanism for Cooperative Information with Application to IV's Safety," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1940–1946.
- [33] J. Müller, T. Griebel, M. Gabb, and M. Buchholz, "Subjective Logic-Based Identification of Markov Chains and Its Application to CAV's Safety," in *Proc. IEEE Connected and Automated Vehicles Symp. (CAVS)*, Honolulu, HI, USA, USA, Sep. 2019, pp. 1–5.
- [34] J. Müller, M. Herrmann, J. Strohbeck, V. Belagiannis, and M. Buchholz, "LACI: Low-effort Automatic Calibration of Infrastructure Sensors," in *Proc. IEEE Int. Conf. Intell. Transp. Syst. (ITSC)*, Auckland, New Zealand, Oct. 2019, pp. 3928–3933.
- [35] J. Müller, J. Strohbeck, M. Herrmann, and M. Buchholz, "Motion Planning for Connected Automated Vehicles at Occluded Intersections with Infrastructure Sensors," *IEEE Trans. Intell. Transp. Syst.*, 2021, under Review.
- [36] D. Nuss, S. Reuter, M. Thom, T. Yuan, G. Krehl, M. Maile, A. Gern, and K. Dietmayer, "A random finite set approach for dynamic occupancy grid maps with real-time application," *The International Journal of Robotics Research*, vol. 37, no. 8, pp. 841–866, 2018.
- [37] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: an open-source Robot Operating System," in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA) Workshop on Open Source Robotics*, Kobe, Japan, May 2009.
- [38] H. Rehborn, M. Koller, Y. Dülger, B. Ünal, M. Maier, and B. Völz, "Energy Consumption of Automated and Non-Automated Vehicles in Various Traffic Scenarios," in *Aachen Colloquium Sustainable Mobility* 29, 2020.
- [39] H. Rehborn, M. Koller, and S. Kaufmann, *Data-Driven Traffic Engineering: Understanding of Traffic and Applications based on Three-Phase Traffic Theory*. Elsevier Science, 2021.
- [40] S. Reuter, B. T. Vo, B. N. Vo, and K. Dietmayer, "The Labeled Multi-Bernoulli Filter," *IEEE Transactions on Signal Processing*, vol. 62, no. 12, pp. 3246–3260, 6 2014. [Online]. Available: <http://ieeexplore.ieee.org/document/6814305/>
- [41] J. Rios-Torres and A. A. Malikopoulos, "A Survey on the Coordination of Connected and Automated Vehicles at Intersections and Merging at Highway On-Ramps," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 5, pp. 1066–1077, 2017.
- [42] SAE International, "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles (SAE J3016)," SAE International, Tech. Rep., 2021.
- [43] J. Strohbeck, V. Belagiannis, J. Müller, M. Schreiber, M. Herrmann, D. Wolf, and M. Buchholz, "Multiple Trajectory Prediction with Deep Temporal and Spatial Convolutional Neural Networks," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1992–1998.
- [44] B. Völz, M. Maier, R.-W. Henn, R. Siegwart, and J. Nieto, "Towards Infrastructure-Supported Planning for Urban Automated Driving," in *Robot: Science and Syst. (RSS), Workshop on Scene and Situation Understanding for Autonomous Driving*, 2019. [Online]. Available: <https://drive.google.com/file/d/1f4ngXPiXILtP71glfNiBNmPFqmSmbxD>
- [45] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal Trajectory Generation for Dynamic Street Scenarios in a Frenet Frame," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Anchorage, AK, USA, May 2010, pp. 987–993.
- [46] H. Winner, S. Hakuli, F. Lotz, and C. Singer, Eds., *Handbook of Driver Assistance Systems*. Springer International Publishing, 2016.
- [47] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A Survey of Autonomous Driving: Common Practices and Emerging Technologies," *IEEE Access*, vol. 8, pp. 58 443–58 469, 2020.
- [48] K. Zhang, Y. Mao, S. Leng, Y. He, and Y. ZHANG, "Mobile-Edge Computing for Vehicular Networks: A Promising Network Paradigm with Predictive Off-Loading," *IEEE Veh. Technol. Mag.*, vol. 12, no. 2, pp. 36–44, 2017.
- [49] J. Ziegler, P. Bender, T. Dang, and C. Stiller, "Trajectory planning for Bertha — A local, continuous method," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Dearborn, MI, USA, Jun. 2014, pp. 450–457.



system identification.



tonomous driving.

**Michael Buchholz** (michael.buchholz@uni-ulm.de) received his Diploma degree in Electrical Engineering and Information Technology as well as his Ph.D. from the faculty of Electrical Engineering and Information Technology at Karlsruhe Institute of Technology, Germany. Since 2009, he is serving as a research group leader and lecturer at the Institute of Measurement, Control and Microtechnology at Ulm University, Germany. His research interests comprise connected automated driving, electric mobility, modelling and control of mechatronic systems, and

**Johannes Müller** (johannes-christian.mueller@uni-ulm.de) received his bachelor's and master's degree in electrical engineering and information technology from Karlsruhe Institute of Technology, Germany in 2013 and 2016, respectively. He is currently a researcher at the Institute of Measurement, Control, and Microtechnology at Ulm University, working towards his Ph.D. degree in the field of motion planning and reliability estimation for connected automated driving.



**Martin Herrmann** (martin.herrmann@uni-ulm.de) received the B.Sc. degree in Media Technology (with first class honors) in 2014 from Ilmenau University of Technology and his M.Sc. degree in Electrical Engineering (with first class honors) in 2017 from Ulm University, both located in Germany. Since 2017 he is a research assistant at the Institute of Measurement, Control and Microtechnology at Ulm University in Germany. His research interests are sensor data fusion, filtering and estimation, signal processing and environment perception for au-



**Jan Strohbeck** (jan.strohbeck@uni-ulm.de) received his Bachelor of Engineering degree in Information Technology and Master of Science degree in Computer Science from Aalen University, Germany. Since 2018, he is working as a researcher at the Institute of Measurement, Control and Microtechnology at Ulm University, Germany. His research focus is motion prediction of traffic participants, specifically at intersections, and the estimation of the prediction's inherent uncertainty.



**Hubert Rehborn** (hubert.rehborn@daimler.com) studied Electrical Engineering and Economics at technical university of Aachen and received his doctorate at University of Hagen. Since 1999 he has been employed at Mercedes-Benz AG Research and Development in traffic research. Today at Mercedes-Benz AG series development in the department of "navigation" he investigates the future of traffic services and is author of the recent book "Data-Driven Traffic Engineering - Understanding of Traffic and Applications Based on Three-Phase Traffic Theory".



**Benjamin Völz** (benjamin.voelz@de.bosch.com) received his Diploma degree from the faculty of Electrical and Computer Engineering at Dresden University of Technology and his Ph.D. from the department of Mechanical and Process engineering at ETH Zurich. In 2017, he joined the Corporate Research of the Robert Bosch GmbH as research engineer focusing on planning for connected urban automated driving. His research interests include situational awareness, prediction, decision making, and planning for both automated vehicles and robots.



**Matthias Maier** (matthias.maier6@de.bosch.com) received his Bachelor of Engineering degree from the faculty of Mechanics and Electronics at Heilbronn University of applied Science (HHN), Germany. Since 2016, he works as software and system developer at the advance development department of Robert Bosch GmbH's business unit Chassis Systems Control. His field of work comprises connected automated driving, teleoperated driving, behavior and motion planning and vehicle control.



**Jonas Paczia** (jonas.paczia@daimler.com) received his bachelor's and master's degree in computer science from Karlsruhe Institute of Technology (KIT), Germany in 2018 and 2021, respectively. He is currently a researcher at Mercedes-Benz AG working toward his Ph.D. degree in the area of energy consumption computation for vehicles.



**Oliver Stein** (oliver.o.stein@daimler.com) received his bachelor's and master's degree in computer science at Esslingen University of Applied Sciences, Germany in 2019 and 2020, respectively. Since 2021 he is working towards his Ph.D. as a researcher at Mercedes-Benz AG in the field of macroscopic traffic analysis with crowd data.



**Rüdiger-Walter Henn** (ruediger.henn@de.bosch.com) holds a Diploma degree and a Ph.D. in science from the physics departments of the Technical University of Darmstadt (TUD) and the Karlsruhe Institute of Technology (KIT), respectively. He became research fellow at the Max-Planck-Institute for Solid State Research, Stuttgart, and joined the Robert Bosch GmbH in 2000 to serve in several management positions in the fields of ADAS and AD. Rüdiger lectures on ADAS at the KIT.