How to Build a Highly Accurate Digital Twin – Intelligent Infrastructure in the Corridor for New Mobility - ACCorD

Laurent **Kloeker**, M.Sc.; Amarin **Kloeker**, M.Sc.; Fabian **Thomsen**, M.Sc.;
Armin **Erraji**, M.Sc.; Univ.-Prof. Dr.-Ing. Lutz **Eckstein**
Institute for Automotive Engineering (ika), RWTH Aachen University,
Aachen, Germany

Contact: laurent.kloeker@ika.rwth-aachen.de

Content

1 Introduction 652
2 Related Work 654
3 Approach 655
4 Implementation 655
4.1 Setting Up the Test Field 655
4.1.1 Construction Measures 655
4.1.2 Hardware Setup 661
4.2 Generation and Utilization of Digital Maps 663
4.3 Traffic Detection via Smart Roadside Infrastructure Sensors 665
4.4 Data Refinement on a Central Computing Server 668
4.4.1 Data Management Concept 668
4.4.2 Description of the Fusion Function 669
4.5 Wireless Data Transfer 672
4.6 Creation of a Digital Twin 673
5 Conclusion and Outlook 674
6 Acknowledgment 675
7 Abbreviations 675
8 References 676

Summary

This paper gives an introduction to the findings of the public-funded research project Corridor for new Mobility Aachen – Düsseldorf (ACCorD). The aim of the project is to build a highly accurate digital twin on three different digital test fields in real road traffic between Aachen and Düsseldorf. The digital test fields and the digital twin are intended for testing and rolling out automated and connected driving functions under real conditions. They consist of an urban, a rural, and a highway section and are equipped with stationary intelligent infrastructure for digital traffic recording. The entire value chain

from the construction of the infrastructure sensor test fields via data processing to the final product of the digital twin is presented in this paper.

1 Introduction

Living Labs, or also called Regulatory Sandboxes, have become increasingly important in the past for the conception, development, and implementation of innovative research approaches. They are by no means limited to mobility concepts of the future, but find a holistic significance in today's research and industrial landscape. The goal of such Living Labs is a new form of cooperation between science and civil society, where the focus is on mutual learning in an experimental environment. Stakeholders from science and industry come together to develop and test scientific and socially robust solutions based on a shared understanding of the problem [1].

With regard to automated and connected mobility, the term "digital test fields" is used primarily. According to the Federal Ministry of Transport and Digital Infrastructure (BMVI), as of September 2020, 14 such digital test fields were already in operation throughout Germany [2]. However, the scope and testing focus of such test fields vary considerably. Nevertheless, two core aspects can be identified. Both the connectivity of vehicles and infrastructure and the collection of traffic data are the focus of these digital test fields. Vehicle-to-everything (V2X) communication marks an essential component of future automated and connected mobility. It not only creates a more comfortable and safer transportation environment but also has great significance for improving traffic efficiency and reducing pollution and accident rates [3]. However, at this stage, the technology is still in the trial phase and has not been fully evaluated for sufficient traffic and information security.

Furthermore, the development and safeguarding of automated driving functions is a strongly data-driven approach [4]. In particular, the collection of large amounts of naturalistic traffic data is a valuable and crucial component in all development workflows. Only through diverse traffic data sets recorded under real conditions a maximum of interaction and scenario complexity can be covered.

The combination of intelligent connectivity and data collection makes it possible to transform a physical system into a digital twin. It contains the properties, states, and behavior of the real-world environment. The advantages of the digital twin are particularly evident in simulation. With regard to digital test fields for automated and connected mobility, for example, this makes it possible to map recorded traffic scenarios and interactions either in real-time or subsequently offline in a virtual environment, depending on availability, and to analyze them there. This offers huge potential in terms of current and future scientific challenges that still need to be overcome on the road to safe and secure automated mobility.

All of the aforementioned aspects are significantly addressed in the research project Corridor for New Mobility Aachen - Düsseldorf (ACCorD) funded by the BMVI [5][6]. The overall objective of the project is to set up and equip individual digital test fields with smart infrastructure sensors for interaction with test vehicles for safe automated and connected driving. In doing so, an integrated development environment will be created by integrating existing test facilities such as the KoMoDnext test field [7] in Düsseldorf or the completed CERMcity test site [8] in Aldenhoven in order to systematically test and safeguard automated vehicles in interaction with connected infrastructure. This is done by a time- and cost-efficient toolchain and methodology, in which simulation, closed test sites as well as public transport test fields are linked in the best possible way. To be able to represent diverse traffic scenarios, the corridor contains a highway section, an urban area, and a rural area.

In detail, the project is divided into the following activities:

- Establishment of a test environment which, with the help of a reference sensor system, records road users and the environment on the three test field sections (urban, rural, highway) with high precision.
- Networking with the traffic infrastructure by means of virtual network control and networked light signal systems.
- Establishment of a central database in which all data collected is processed and stored and can be used for further research and development activities and for simulations.
- Design and implementation of a digital twin of the test field as a virtual image of the three test field sections for the execution of tests in simulation.
- Use of the test environment for the validation of automated and connected test vehicles, which are used in the corridor for the further development of automated driving functions.

In addition to the Institute for Automotive Engineering (ika) - RWTH Aachen University as the project coordinator, the project consortium consists of partners from research, industry, and regional authorities and municipalities. These include the Chair and Institute of Highway Engineering (ISAC) - RWTH Aachen University, the companies e.GO MOOVE GmbH, Ford-Werke GmbH, PTV Planung Transport Verkehr AG, Vodafone GmbH and ZF Friedrichshafen AG, as well as the North Rhine-Westphalian state road construction authority and the city of Aachen. The consortium's broad spectrum of expertise means that all research topics such as automated driving, V2X communication, infrastructure measures, digital twins, and sensor validation can be researched and tested.

The project runs from 01.01.2020 to 31.12.2021 with a total project volume of around 11.11 million euros, of which the BMVI's funding share is around 9.57 million euros.

In the following chapters, we provide an overview of the entire value chain of the ACCorD research project. We explain all intermediate steps from the setup of the digital infrastructure sensor test fields via data processing to the final product of the digital twin.

2 Related Work

As outlined earlier, there are a large number of digital test fields in Germany. In the following, however, we limit to the test fields of the research projects Providentia++ [9][10] and Diginet-PS [11][12] as well as to the Test Bed Lower Saxony [13][14] due to their size and the parallels to the research project ACCorD. Additionally, we address the HDV-Mess research project, which deals with mobile rather than stationary Intelligent Transport Systems Stations (ITS-Ss), yet uses the same infrastructure sensor setup as the ACCorD research project [15][16][17].

In the first and already completed phase, as part of the Providentia project, camera and radar sensors were mounted on two gantries on the A9 highway in Munich. The installed technology is capable of recognizing vehicle types and classes and communicating the data via 5G to create a digital twin of the scene [10]. In the follow-up project Providentia++, the next step is to refine the deployed technology, extend the infrastructure into the "urban space" and develop so-called value-added services. New transmission and sensor masts are to be set up on the edge of a highway feeder road right into inhabited areas [9].

The Diginet-PS research project addresses the three core topics of intelligent vehicle, intelligent road infrastructure, and cloud in an urban environment in Berlin. The goal is to connect three independent entities (vehicle, road infrastructure, and cloud) that are autonomous in their decision-making and to exchange necessary information among them. This information includes both the current situation on the road and forecasts about future situations. The digital test field extends over a length of 3.7 kilometers and covers three lanes in each direction. Different types of sensors allow traffic counts to be carried out and road and environmental conditions to be collected in real-time [11][12].

The Test Bed Lower Saxony for automated and connected mobility, built by the German Aerospace Center (DLR), includes both highway sections and rural road sections. It also integrates the established routes of the Application Platform Intelligent Mobility (AIM) [18], which is in operation in downtown Braunschweig. In total, the Test Bed Lower Saxony will cover more than 280 kilometers of roadway with a wide range of technical equipment. The main equipment used will be roadside recording systems. They record the traffic situation of the road sections and are able to communicate the collected information via V2X to selected receivers [13][14].

In contrast to the previous research projects, the HDV-Mess research project pursues the goal of highly accurate digital traffic recording using mobile and modular measurement setups. Four mobile ITS-Ss, as well as a mobile receiving server, are used. Each measuring station is equipped, among other things, with two LiDARs and two video cameras with which the surrounding traffic is recorded. The recorded raw sensor data is used to detect road users with the help of artificial neural networks and to extract their trajectories for later development purposes. The fully self-sufficient systems allow traffic data to be recorded under real conditions on any road section over a period of several weeks. Following such a measurement, the collected sensor data is transferred to a central computing server for further research purposes [15][16][17].

3 Approach

A lot of complex preparatory work is needed to set up a digital twin. First, technical, structural, and regulatory framework conditions must be created to install and permanently operate intelligent infrastructure sensors in road traffic. In the next step, algorithms must be developed that stabilize, globally register and process the sensor data into object data with the lowest possible latency. This object data is then suitable for wireless transmission to a central computing server, which in turn is capable of refining the data and forwarding it to either to databases, simulation environments or connected and automated vehicles (CAVs).

In the following chapters, we describe in detail the implementation of all these required intermediate steps to achieve the overall goal of a real-time capable highly accurate digital twin.

4 Implementation

The implementation of the digital test fields is an extensive process. In the first step, the hardware and software framework conditions of all components of the entire value chain must be defined. Planning and implementation of construction measures form the basis for *setting up the test fields*, in which the infrastructure sensors are installed at the roadside of the three test fields. Furthermore, the creation and utilization of digital maps is a necessary component for traffic detection via smart roadside infrastructure sensors. The collected traffic data will then undergo data refinement on a central computing server. Through wireless data transfer, the refined measurement data can finally be used for the creation of a digital twin.

4.1 Setting Up the Test Field

As a basis for setting up a digital twin, the structural and hardware requirements must first be met. These will be explained in detail in the following two chapters.

4.1.1 Construction Measures

When determining the location of the test field sections, it was initially crucial that all three domains, urban, rural, and highway, are covered. Only by holistically covering these three domains can the Corridor for New Mobility meet the requirement of being a versatile and diverse digital test field in real traffic. In addition, each of the three test field sections should contain as many interactions between road users as possible. For this reason, requirements were set in advance for the test field sections, which are listed in Fig. 1.

Furthermore, it must be possible to set up the corresponding infrastructure sensors at the selected test field sections. In order to keep occlusions of road users as low as possible, the prerequisite must be met that the infrastructure sensor technology is installed at a height of several meters. This is only possible if existing infrastructure, such as lighting masts, can be used or if new masts are erected at the roadside. Particularly in rural areas and on the highway, test field sections must therefore be selected whose verge and embankment are wide enough for such a project. In urban areas, the aspect of verges and embankments is not applicable due to the already dense development. Taking into account all the aforementioned requirements, three test field sections were therefore found in the area between Aachen and Düsseldorf, each of which fulfilled all the criteria.

	Urban	**Rural**	**Highway**
Section elements	• Intersection • Roundabout • Bus stop	• Intersection • Side strip	• Slip road
Road user classes	• Cars • Trucks • Busses • Motorcycles • Bicycles • Electric scooters • Pedestrians	• Cars • Trucks • Busses • Agricultural vehicles • Motorcycles	• Cars • Trucks • Busses • Motorcycles

Fig. 1 Requirements for the three test field sections.

The urban test field section is located in Aachen on the Melaten Campus. It consists of Seffenter Weg, Campus Boulevard and Forckenbeckstrasse. In total, it forms a circuit of 2.4 kilometers. As can be seen in Fig. 2, it contains the required section elements of intersection, traffic circle, and bus stops. Additionally, all required road user classes move along this test area section. Due to the location of research institutes of the RWTH Aachen University and companies, there is a lot of traffic using private and public transport as well as shared micro-mobility services. The prescribed speed limit is 50 kilometers per hour.

To be able to cover the entire test field section, the existing lighting infrastructure is used. The lighting masts on the inside sidewalk side are ideally suited for this purpose with a height of eight meters. A total of 46 measurement locations were defined. First, every second lantern on the circuit was selected. Subsequently, additional surrounding lanterns were selected at the southern intersection as well as at the northern rotary in order to completely cover these highly interactive sections in particular. Finally, an additional measurement location was added at each of the seven bus stops on the inner side of the street in order to fully cover in particular safety-relevant information about pedestrians who, for example, want to change roadsides while hidden behind a bus. All 46 measurement locations are also shown in Fig. 2.

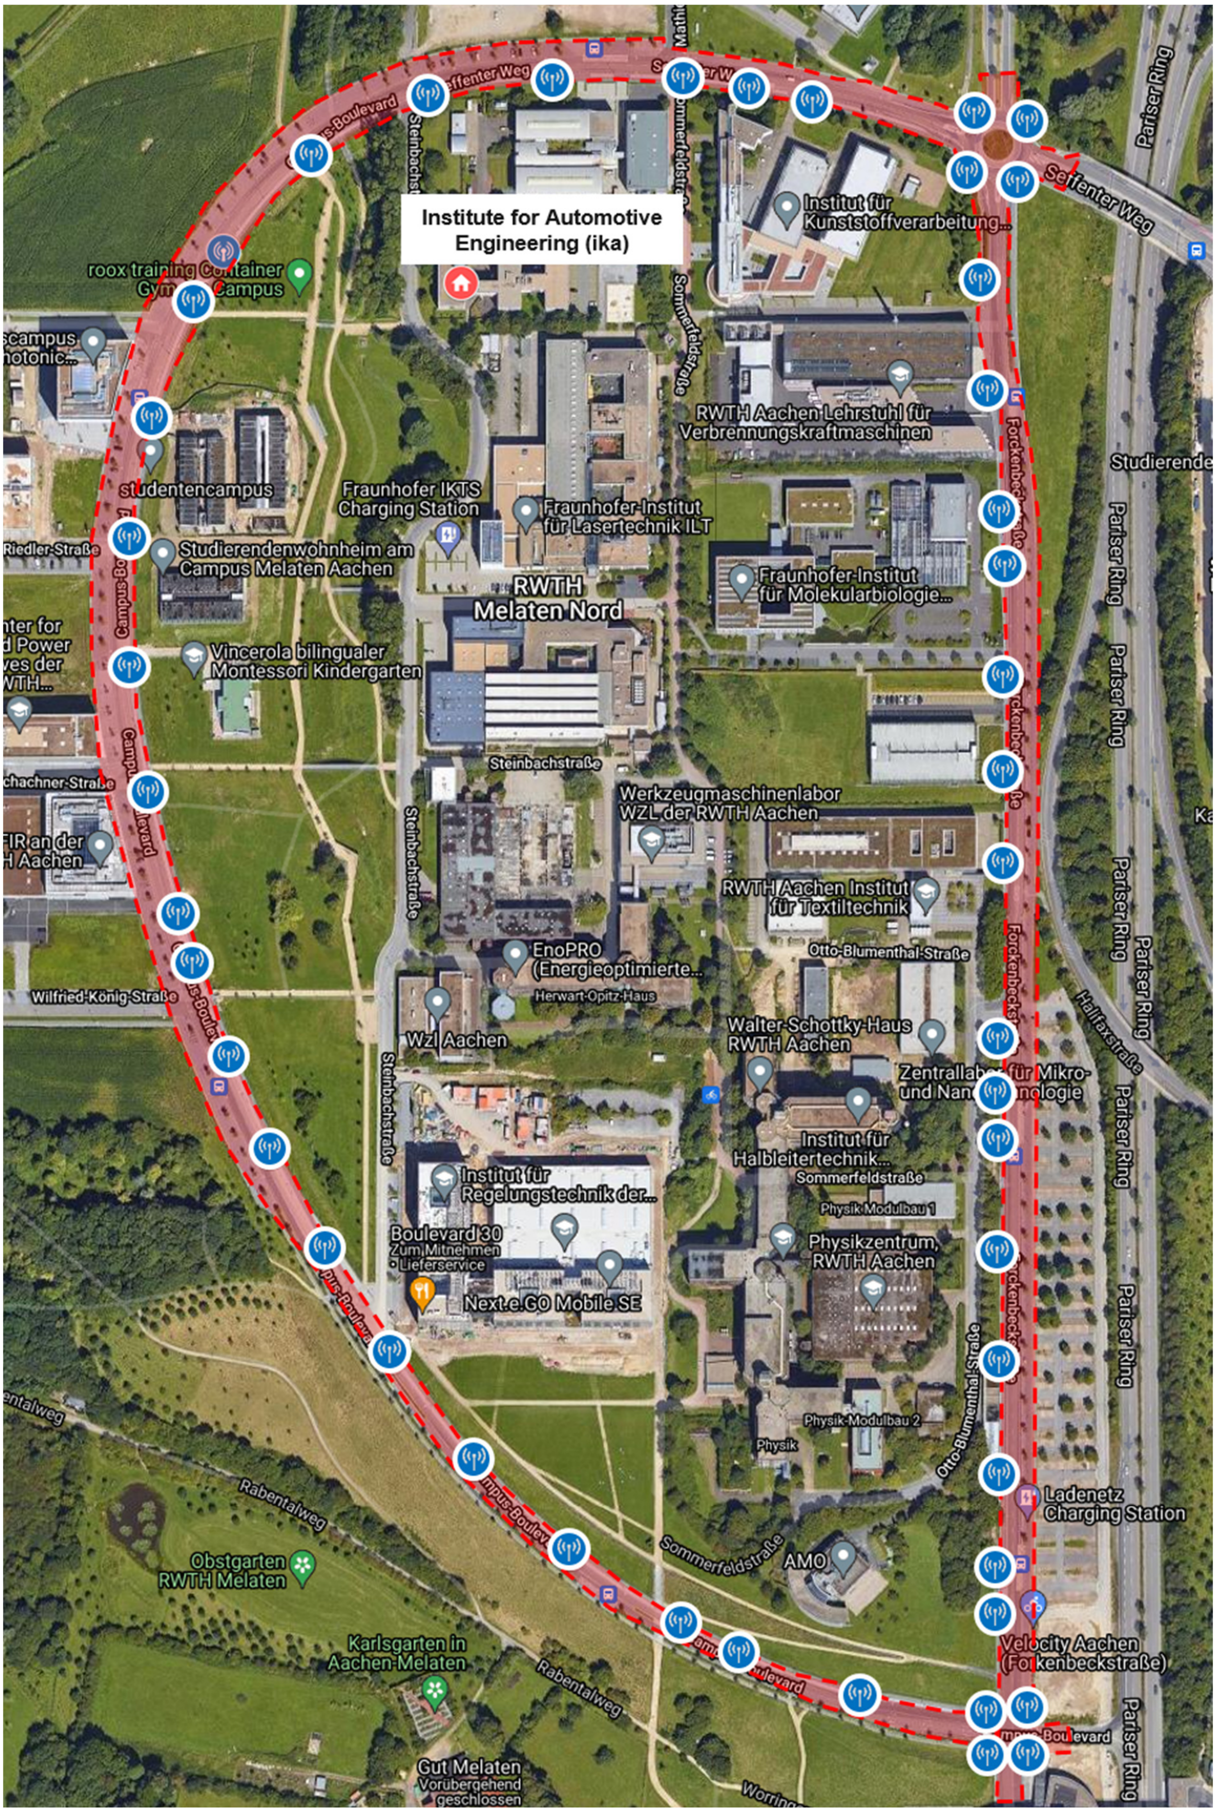

Fig. 2 Illustration of the urban test field section on the Campus Melaten in Aachen. Background image source: [19].

The rural test field section is located on the B56 near Dürboslar at Aldenhoven in the immediate vicinity of the Aldenhoven Testing Center (ATC). As can be seen in Fig. 3, it extends over a straight length of one kilometer and includes an intersection, a junction of a service road, and wide shoulder strips on both sides of the road. In addition, on the southbound side of the roadway, both the verge and the embankment are designed to be sufficiently wide for the construction of new masts, since, unlike the test field section on the Campus Melaten, no lighting masts are placed at the edge of the road. Due to the surrounding fields, many agricultural vehicles drive on this test field section and, because of their low speed, use the wide shoulder to give following traffic the opportunity to overtake with sufficient safety distance. The prescribed speed limit in the intersection area is 70 kilometers per hour and 100 kilometers per hour on the rest of the test field section. No sidewalks run along the roadside, so pedestrians tend not to be included among the road users.

Fig. 3 Illustration of the rural test field section on the B56 near Aldenhoven. Background image source: [19].

The new masts to be erected will be set up at a distance of 100 meters from each other so that a total of eleven measurement sites will be available. These are also shown in Fig. 3. Due to safety regulations, a new continuous guardrail will also be installed at the roadside to prevent possible collisions with the masts in the event of traffic accidents. In order to create the same conditions as at the Campus Melaten test field section, masts with a height of six meters will be installed so that the infrastructure sensors can be installed at a sufficient height.

The third test field section is located on highway A44 at the Jackerath interchange. Similar to the rural test field section, it extends over a length of one kilometer, whereby the roadway is slightly curved. Due to the existing lane widths, only one driving direction of this test field section can be covered, which is why the southern lane with driving direction Düsseldorf was selected. According to the requirements of Fig. 1, the test field section contains two highway slip roads in immediate succession. The first highway slip road is single-lane and merges with a two-lane roadway. The second highway slip road is two lanes and also merges with a two-lane roadway. At the end of the acceleration lane, the test field section merges into a three-lane roadway. As can be seen in Fig. 4, the entire length of both highway slip roads is fully covered. Due to the two highway slip roads, the road users of the test field section are forced to change lanes or interact with each other. It is from these two aspects that valuable insights can be gained with the measurement data collected in the future. Along the entire section, both the verge and the embankment are also sufficiently wide to allow new masts to be erected at the roadside. The construction of new guardrails is in turn not necessary, as the existing guard rails are already sufficient for this project. The prescribed speed limit on the highway slip roads is 80 kilometers per hour. On the main lanes, it is unlimited.

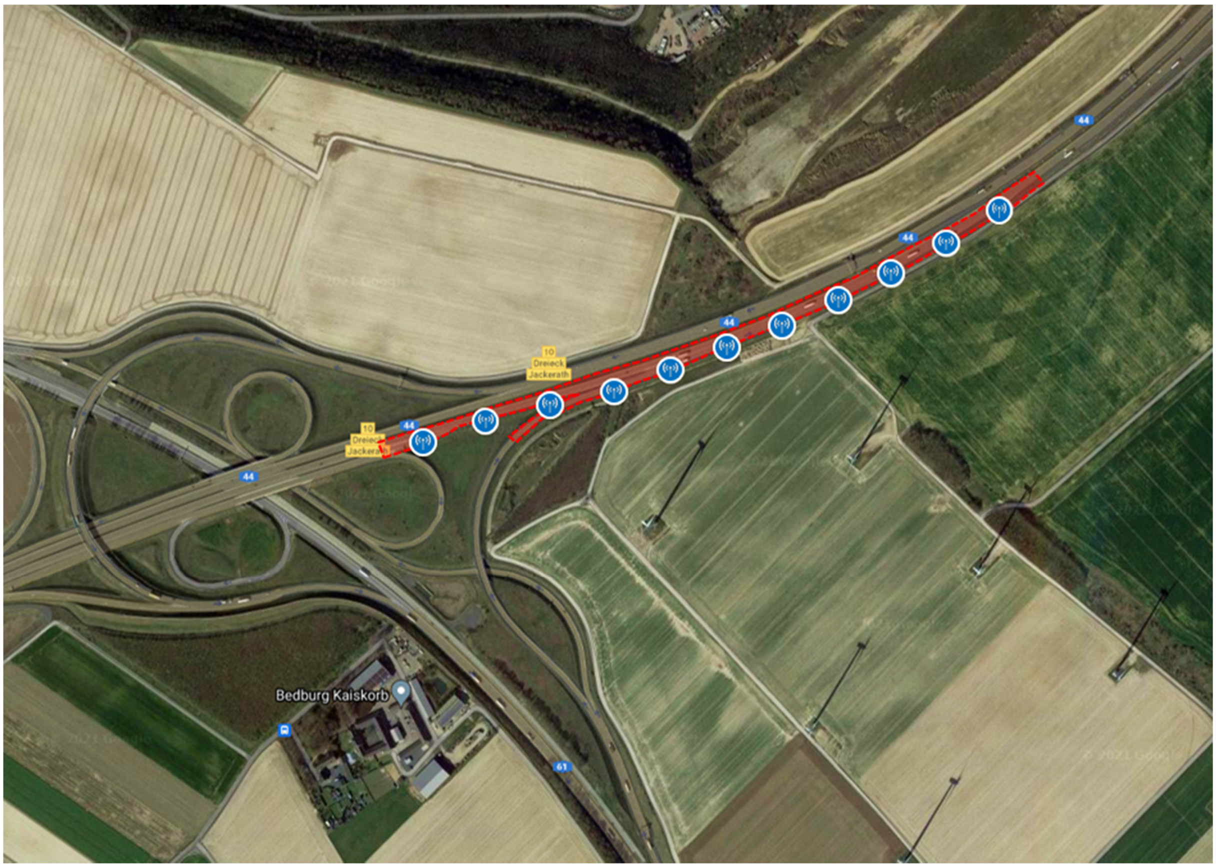

Fig. 4 Illustration of the highway test field section on the A44 at the Jackerath interchange. Background image source: [19].

On the highway test field section, the masts are also erected at a distance of 100 meters from each other at the edge of the road. As can be seen in Fig. 4, the first three measurement sites are located within the highway interchange and the remaining eight measurement sites are located at the southern edge of the roadway. The new masts to be erected will also have a height of six meters so that there is an equal mounting height of the infrastructure sensors on all three test field sections.

The issue of energy and data supply at all three test sites proved to be a fundamental challenge. For efficiency reasons with regard to raw sensor data processing, a connection of all measurement sites to a fiber-optic network was initially planned. This could forward the raw data to an edge computing server and process it collectively in real-time. However, on the test field section Campus Melaten, a connection of the lighting masts to a fiber optic cable network was unfortunately not possible due to structural reasons. This ultimately meant that the raw data had to be processed locally at each measurement site individually. Local raw data processing in turn requires powerful measuring computers for each measuring site, which have a permanent power supply and increased energy requirements. Thus, all 46 lighting masts on the Campus Melaten test field section as well as the associated distribution nodes were upgraded accordingly in order to implement the requirements. All components of the infrastructure sensors and the measurement computer can be supplied with sufficient power via a mast socket at the appropriate height.

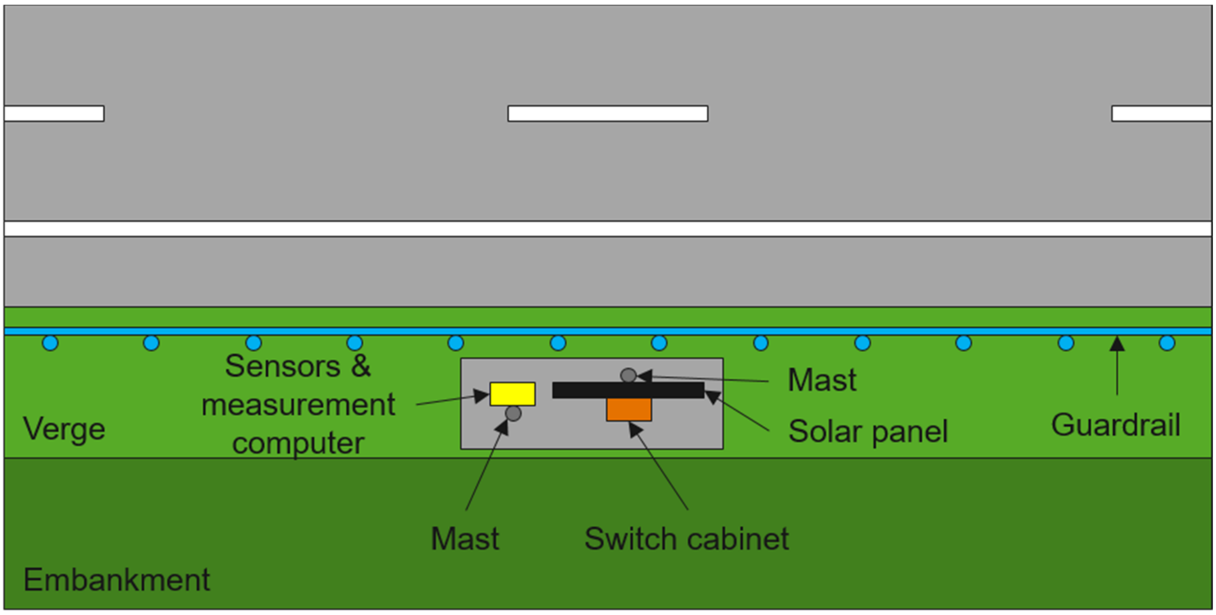

Fig. 5 Schematic illustration of the setups of infrastructure sensors and local fully self-sufficient power supply at test field section B56.

On the test field sections B56 and A44, it was examined at the beginning whether, on the one hand, an existing energy supply network can be accessed, or whether it is possible to lay a fiber optic cable network. However, since both aspects also had to be ruled out for construction and regulatory reasons, a local and fully self-sufficient energy supply was chosen. As on the test field section Campus Melaten, the new energy supply provides sufficient electrical energy for the components of the infrastructure sensors as well as for the measurement computer for local raw data processing at each measurement site. Large-scale solar panels are used in combination with buffer batteries. In addition to the mast for the infrastructure sensor system, each measurement site of the test field sections B56 and A44 will receive another neighboring mast on which the solar panels will be installed. This reduces vibrations in the sensor mast that would otherwise be induced by the solar panels, which are susceptible to wind, and would affect the measurement data quality. The buffer batteries are housed in a control cabinet at the base of the mast. A cable from the control cabinet, to the sensor mast, provides the power supply. A schematic diagram of a measurement site of the test field section B56 with the sensor mast and the power supply mast is shown in Fig. 5.

4.1.2 Hardware Setup

The hardware setup of our smart roadside infrastructure sensors is divided into two parts, a sensor module and a control cabinet which contains all the remaining components required for operating a measurement station. The sensor module is mounted at a height of about six meters either on a lighting mast or a newly erected mast to get a view of the traffic from a perspective as high as possible. The corresponding control cabinet is mounted a few meters lower in order to keep the center of gravity of the entire setup as low as possible and thus minimize unwanted oscillations of the mast. A schematic of the two modules is shown in Fig. 6.

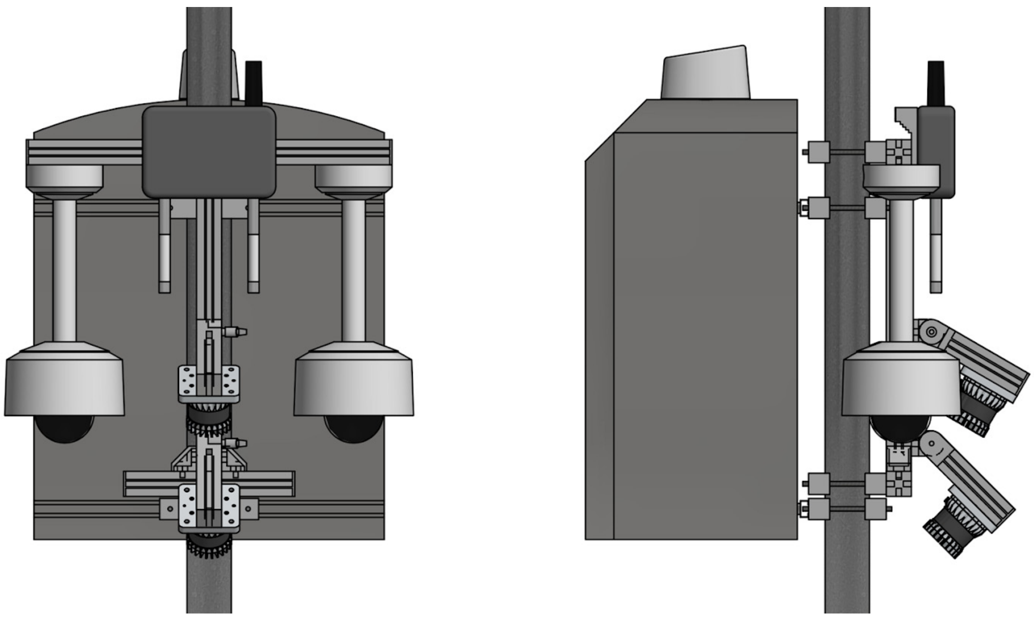

Fig. 6 Computer-aided design (CAD) drawing of a measurement station. For visualization purposes, the sensor module and control cabinet are mounted at the same height.

The sensor modules consist of two LiDARs, two cameras, and a V2X roadside unit (RSU). The LiDARs are used to capture the immediate environment of a measurement station. For this purpose, we use 360-degree rotating LiDARs of the type OS1-128 Gen2 from the company Ouster. These sensors offer a vertical angular resolution of 0.35 degrees with a vertical field of view of 45 degrees and thus generate a very dense point cloud which serves as a basis for highly accurate object detection. The angled arrangement of the two LiDARs allows us to cover the entire width of the traffic corridors. In overlapping areas, the density of the point cloud and thus the accuracy of object detection is additionally increased.

The cameras are mainly used to support the LiDARs to provide additional information when the LiDARs cannot fully detect or classify an object due to occlusion. For this purpose, we use pan, tilt, roll, zoom (PTRZ) dome cameras of the type FLEXIDOME IP starlight 8000i from Bosch. This allows us an individual and optimal alignment of the cameras for each measuring station. Here, the cameras are not aligned to the viewing area of the LiDARs of the own measuring station, but one camera each is aligned to a neighboring measuring station. In this way, the area around a measuring station is to be captured from as many perspectives as possible, thus minimizing occlusions. The high resolution of the cameras of 3840x2160 pixels allows us to detect objects with high precision also in the camera images. A schematic representation of the sensors' fields of view (FOVs) can be found in Fig. 7.

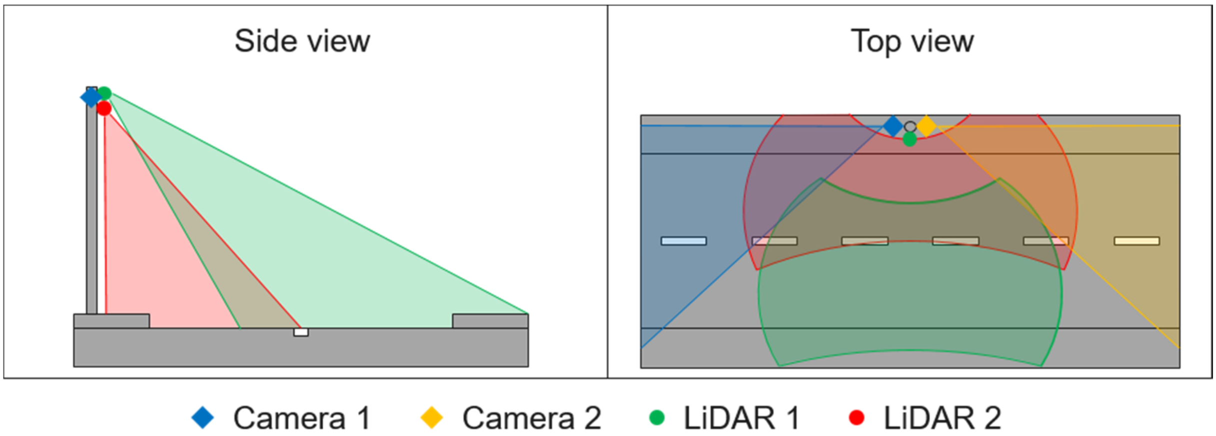

Fig. 7 Schematic representation of the sensors' FOVs of an ITS-S. Note: The scales shown in the diagram on the right are not true to reality.

A measuring station is completed by the control cabinet, in which a high-performance computing unit and a multi-band antenna for data communication are installed. The hardware of the computing unit is designed in such a way, that at full load together with the sensor module our power limit of 475 Watts per station is not exceeded, but still as many central processing unit (CPU) and graphics processing unit (GPU) resources as possible are available for direct data processing. According to the state-of-the-art at the time of the construction of the measuring stations, an AMD Ryzen 7 3700X CPU and an Nvidia Quadro RTX6000 GPU are installed. The computing system is completed by 64 gigabytes random-access memory (RAM) with automatic bit error correction (ECC) to prevent data corruption during operation.

For communication purposes, the computing unit has enough Power-over-Ethernet (PoE) network ports to connect all external devices and leave room for possible future expansions. On the other hand, various wireless modules are installed, which allow connections to other devices via the Navilock NL-400 multi-band antenna. Most important is the 4G module, which is used to connect to our central server via a cellular network. For this, we use a Sierra Wireless EM7565 module, which is directly integrated into the system and connected to the multi-band antenna. An upgrade to 5G for testing the new mobile radio standard is possible and planned for the future. A GNSS receiver is also directly integrated into the 4G module, which serves as a source for precise time information. This ensures that all measuring stations are temporally synchronized. In the event of a cellular connection failure, a WiFi module is also integrated as a backup. It provides an access point for each measuring station, which allows local wireless maintenance of the system. In case of software or network problems, this avoids time-consuming maintenance directly on the hardware at a height of six meters. For local ITS-G5-based V2X communication, one MK5 RSU from Cohda Wireless is also installed for each measuring station. In parallel to cellular-V2X (C-V2X) via 4G, it represents the second communication channel for object messages.

4.2 Generation and Utilization of Digital Maps

Within the context of both intelligent infrastructure and highly automated vehicles, digital maps have gained importance over the last few years. In contrast to publicly available maps, they offer information about an environment with higher accuracy and level of detail, which is why they are also referred to as high definition maps. This section gives an overview of different specialized formats that are used in the ACCorD project together with their respective applications. The specific maps were mostly created by Atlatec GmbH, a company that specialized in the recording of high-definition maps.

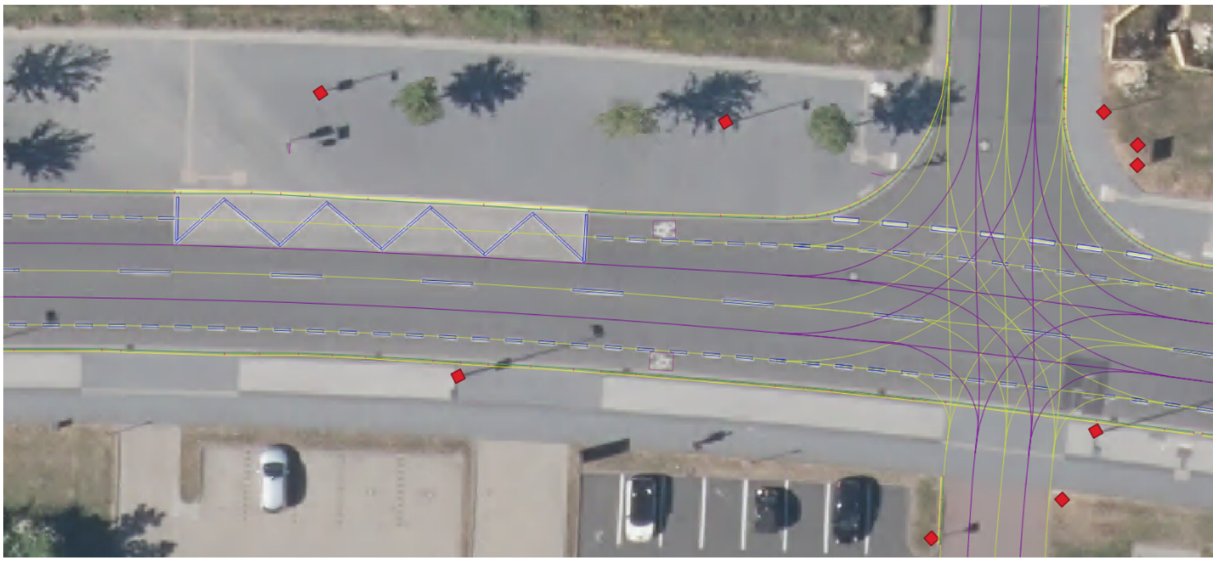

Fig. 8 A Shapefile projected onto an orthophoto of the urban test field.

Fig. 8 shows an extract of a Shapefile [20]. This format was introduced by the *Environmental Systems Research Institute* (ESRI) for the storage of general vectorized geodata. Its focus lies on the exact representation of geometries as two- or in this case three-dimensional polygons in a geo-referenced coordinate system. Each shape within the file can be linked to one or more categories and additional metadata of which the type is the most important one. The projection of all outlines onto an orthophoto of the urban test field shows the different types of structures that are stored in the file: Lane markings and other road markings, e.g. the diagonal lines indicating the space for a bus stop, exactly match up with the visible shapes. For more complex road paintings like bike lane indicators, their course outline is given. Borders and centerlines of all tracks are extrapolated from the information on the road. Finally, outlines of curbs and the positions of all masts around the street, especially road signs and street lamps, are indicated on the map. The latter are of particular interest in the ACCorD project providing precise locations of the ITS-Ss. Overall, the Shapefile can mostly be used for sensor localization. With its exact representation of geometries that are visible for cameras, these are able to determine their own pose within the test field. Moreover, by mapping pixels in their image onto three-dimensional points in the Shapefile, the cameras can be calibrated for the compensation of lens distortion.

While applicable for cameras, the registration of LiDAR sensors is more usefully performed using a georeferenced point cloud. The recording of such is usually realized by a LiDAR sensor mounted to a vehicle equipped with a global navigation satellite system (GNSS) and possibly an inertial measurement unit (IMU). Using the latter two, the recorded raw points are combined into a single map using simultaneous localization and mapping (SLAM) algorithms.

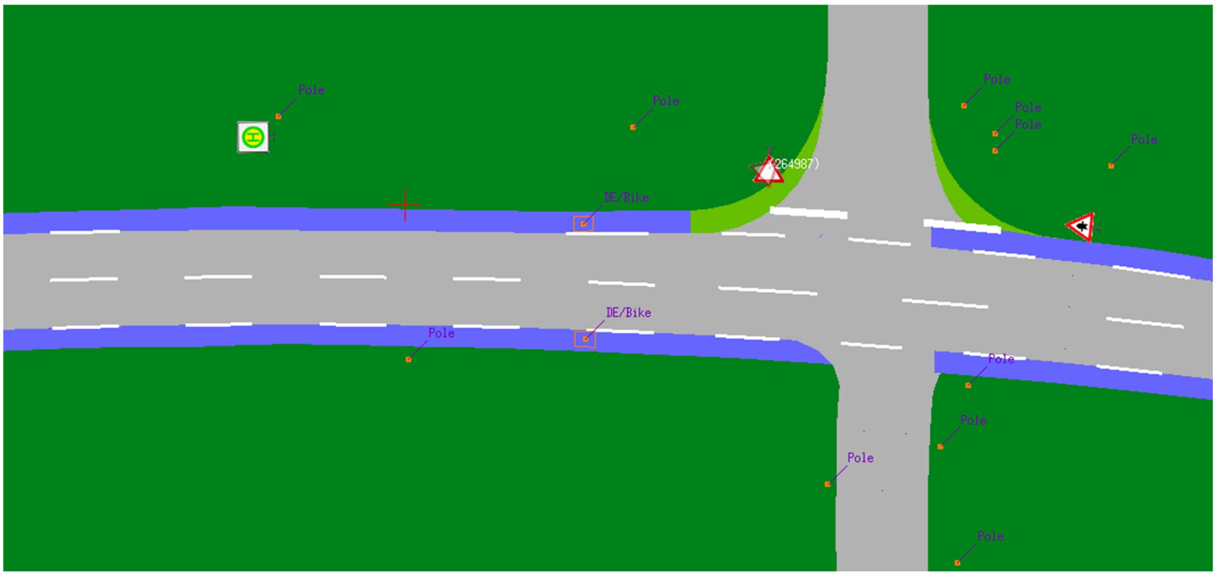

Fig. 9 An OpenDRIVE file of the urban test field.

The focus of other formats mainly lies in semantic relations between objects and their implications for traffic. Therefore, the lanes are not represented as individual shapes but as a combined road with additional information attached to them. This can be for example the velocity allowed, restrictions for vehicle types e.g. on bike lanes, all possible successor lanes on intersections, or rules for the right of way for crossing lanes. Likewise, lane markings or signs are stored as logical elements rather than their exact geometries. For example, a dashed line in these formats does not comprehend the position and length of each dash but instead their meaning, which is that vehicles are allowed to cross the line from both directions. Multiple formats of this type currently exist and are used in different software applications. The OpenDRIVE standard, originally released by *VIRES Simulationtechnologie GmbH* and now maintained by the *association for standardization of automation and measuring systems* (ASAM), is mostly used by simulations like Virtual Test Drive [21]. This application is able to replicate scenarios with multiple vehicles and generate raw data of sensors like cameras, LiDARs and Radars. Fig. 9 shows the same part of the test field as Fig. 8 in an OpenDRIVE format clearly manifesting the most important differences. IPG CarMaker, another simulation software focused on the virtual testing of automated driving functions, uses the Road5 format to which OpenDRIVE maps are directly convertible [22]. Finally, the Lanelet2 format introduced by *FZI Forschungszentrum Informatik* is based on OpenStreetMap representations [23]. It can therefore be edited and viewed with the same tools while delivering libraries for the usage in various programming languages and frameworks, especially the Robot Operating System (ROS). This framework providing tools for the processing of sensor data of all types as well as communication within distributed systems is widely used in the prototypical development of automated driving functions. For this reason, Lanelet2 maps are widely used within this field of research and play a central role in the testing of driving functions within a digital twin.

4.3 Traffic Detection via Smart Roadside Infrastructure Sensors

Thanks to a large amount of computing power, we are able to perform many data processing steps directly on the measuring station and can thus drastically reduce the amount of data that has to be sent to the central server. Only in this way can we achieve realistic data rates for the bandwidth of the mobile network, while at the same time minimizing latency. When processing the data, all sensors are considered individually. However, due to the overlapping field of views, the data of the LiDARs can be merged in the course of the processing pipeline, which can be seen in Fig. 10.

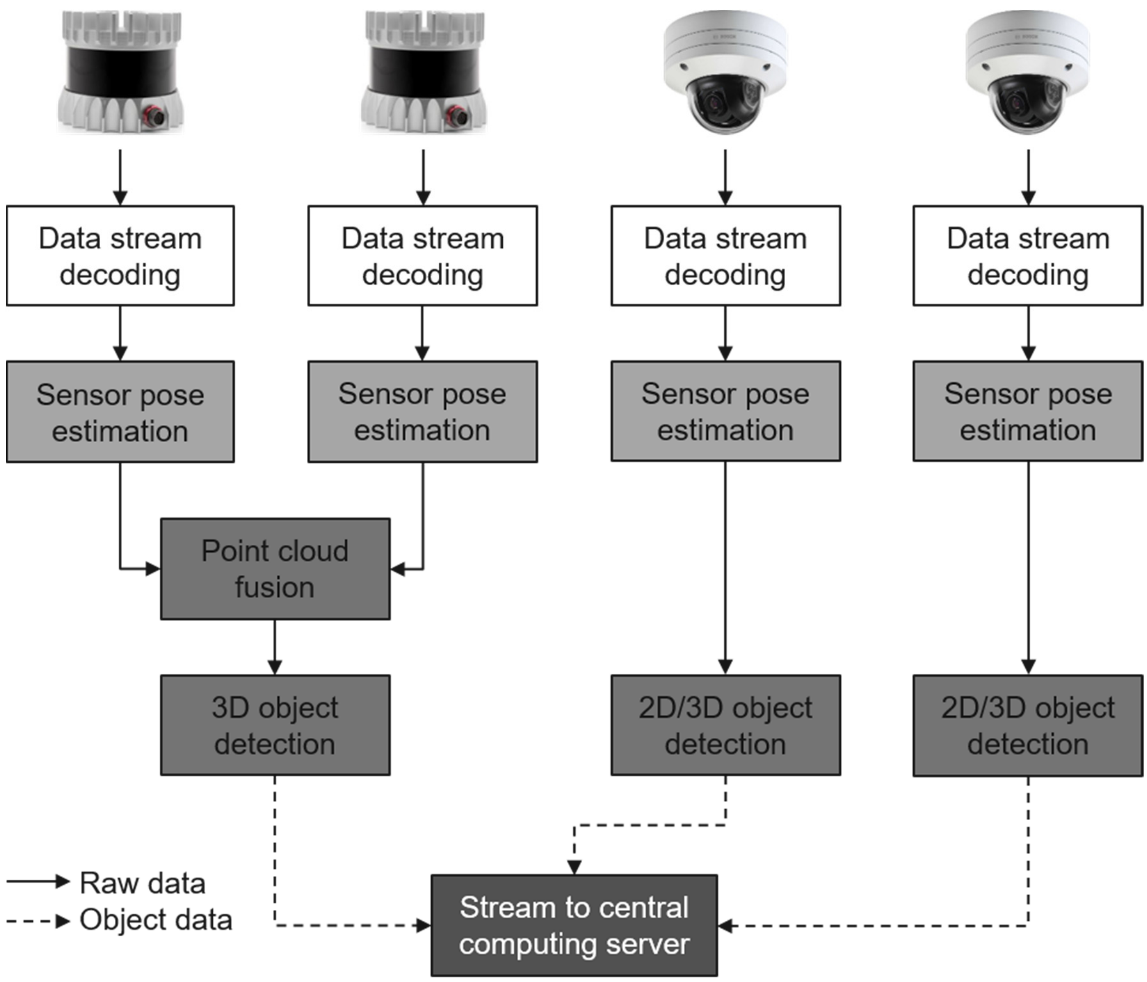

Fig. 10 Data processing pipeline of our roadside infrastructure sensor stations. Sensor image sources: [24][25].

The first step in processing the LiDAR data is decoding the raw sensor data. The LiDAR data is sent via Ethernet using a proprietary codec and decoded at the measurement station using a custom driver. Subsequently, the LiDAR data is available as a point cloud in a Cartesian coordinate system. Next, the poses of the two LiDARs are determined by registering the respective point clouds onto the digital maps described in section 4.2. Since the global pose of the LiDARs is unknown at the beginning of a measurement, it is first determined with the help of the IMU of the sensor and the GNSS of the measurement station by a computationally intensive point cloud registration. For this purpose, both feature-based and point-based methods are used, whereby a highly accurate registration can be achieved. In the further course of a measurement, the pose must be permanently updated, since it constantly fluctuates due to environmental influences such as wind and vibrations. Thanks to the initial registration, however, a purely point-based procedure, which can be carried out efficiently in real-time and does not introduce any further significant errors to the initial registration, is sufficient for this. Once the pose of the two LiDARs is determined, their point clouds can be transformed into the global coordinate system defined by the digital maps. This also means that by simply concatenating the two point clouds, they are directly fused. Thus, it is sufficient to perform the subsequent object detection only once on the fused point cloud. At the same time, the higher resolution of the point cloud in the overlapping regions of the two LiDARs increases the detection accuracy in these regions. A state-of-the-art neural network is used for object detection, which provides a list of 3D objects in the point cloud as output. The network is trained to detect all expected dynamic objects such as pedestrians, animals, bicyclists, cars, etc. An example of object detection on a fused point cloud is shown in Fig. 11.

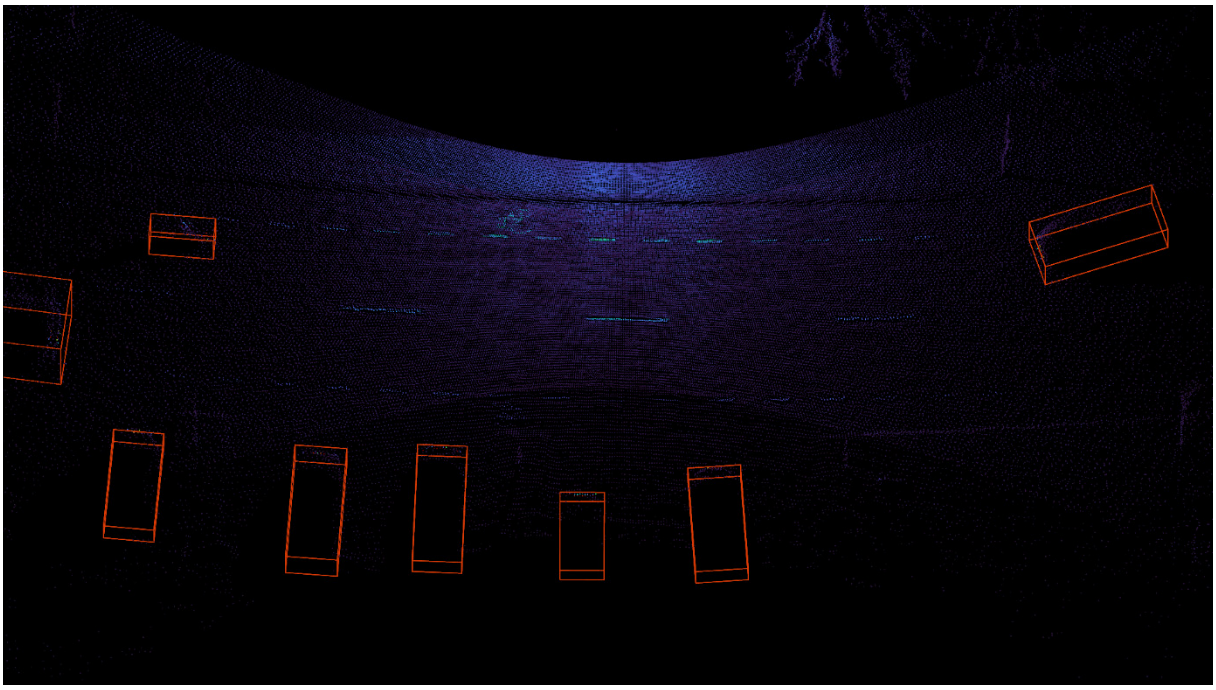

Fig. 11 Visualization of a fused point cloud recorded by an ITS-S with bounding boxes of detected objects highlighted in red.

Just like with LiDARs, the first step in the camera data processing pipeline is the decoding of the incoming stream. The cameras send their video stream with a resolution of 3840x2160 pixels encoded using H.264 compression, which can be read with the x264 decoder.

The next step is the calibration of the cameras. This consists of two substeps: First, intrinsic calibration is used to remove distortions caused by the lens itself (the relevant parameters here are the focal length and principal point) and inaccuracies in the manufacturing process (radial and tangential distortion) of cameras. The intrinsic parameters therefore only have to be determined once, which is classically done with a calibration pattern, such as a checkerboard image filmed from different perspectives. However, this approach is not practical with a large number of cameras, so instead a matching between camera images and the digital map, more precisely the polygons in the Shapefile, takes place. Fig. 8 shows the road markings in this, Fig. 12 demonstrates how they look in camera images. Using classical image processing algorithms, these are detected, filtered according to their quality, and finally mapped to those in the digital map. From several such images, the intrinsic parameters can be calculated according to the checkerboard method. The second substep is the extrinsic calibration, also called registration, which is performed permanently in operation, analogous to LiDARs. Using the same features in the images and the map, the pose of the camera is determined and this information is used to stabilize the video stream.

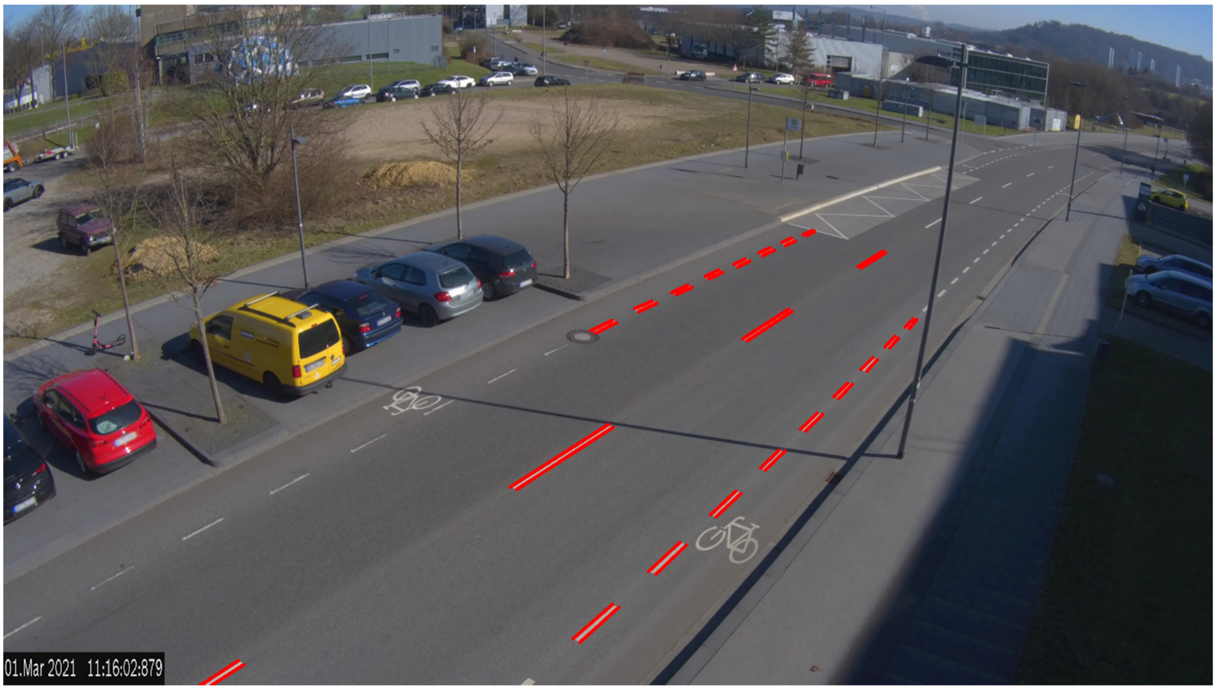

Fig. 12 Detected lane markings in a camera image.

Subsequently, neural networks are used to detect objects in the images. Initially, the aim is to achieve two-dimensional segmentation, which shall be extended to three-dimensional object detection in the future. This allows the same data to be recorded as with LiDARs, although we expect that in particular the classification of objects will be carried out with greater accuracy. For example, trucks and buses can hardly be distinguished in a point cloud, but clearly based on the camera image.

With the data of all sensors processed, the generated object lists are sent to the server. Their further processing is described in the following chapter.

4.4 Data Refinement on a Central Computing Server

Centerpiece of the test fields is the central computing server, on which the object lists of the measuring stations are received and processed. Its purpose is to process the data from all ITS-Ss and to generate as much information as possible from the object lists it receives. The following two sections outline the data management concept and provide an in-depth look at our object fusion algorithm.

4.4.1 Data Management Concept

An overview of our data management concept can be seen in Fig. 13. The system is designed to be completely scalable and each node can in theory run on a different server. This allows the computing power to be variably adapted to the workload. However, for now the system is implemented by a single computer with two AMD Epyc Rome 7742 CPUs, 1 terabytes RAM, four Nvidia HGX A100 GPUs and around 200 terabytes SSD storage in a redundant array of independent disks (RAID) of type RAID 10. An extension to multiple machines is planned in order to eliminate the single machine as a single point of failure.

The complete system is based on the open source Apache NiFi project, which is ideal for automating data flows between software systems. In this system, individual data processes can be represented as scalable nodes and their connections can be configured as desired. The measurement stations are not directly integrated into the NiFi system, but have their own stripped-down NiFi version called MiNiFi running on each unit. These MiNiFi nodes send the object lists generated on the ITS-Ss to the central NiFi system and are received there by a *receiving node*, which forwards the data after analysis for correctness and completeness to a *fusion input node*. This node checks from which test field the data originates and then assigns the data to the corresponding fusion function. There is a separate fusion function for each test field. The fusion function generates fused objects for each time step from the input data, as well as complete trajectories as soon as a detected object has left a test field. Further details about the object fusion are explained in section 4.4.2. If the output data of the fusion function is a fused object, it is forwarded by the *fusion output node* to the CV2X interface and to the live visualization. This ensures that these two modules always have the most up-to-date data available. If the Fusion output data is a completed trajectory, it is stored in the trajectory database. When this happens, a scenario extraction for this trajectory is triggered at the same time. For this purpose, the *scenario input node* collects all necessary data from the trajectory database and the database for the static environment and forwards them to a scenario extraction function. Any number of scenario extraction functions can be executed in parallel, which allows it to perform complex calculations. The results of the scenario extraction are then stored in a scenario database. As the development of the scenario extraction is not within the scope of the ACCorD project, we don't provide further details in this paper. The previously mentioned database for the static environment contains all digital maps of all test fields. If the maps in the database are updated, the updates are also automatically transmitted to all ITS-Ss. To work on the generated data, a Python kernel was set up, which has access to all databases. On top of this a Jupyter Notebook Hub is running, which allows easy data handling from external devices. Another advantage of this is that the data does not have to be downloaded to work with it. Of course, it is also possible to directly download all data for further use-cases. For this purpose a distribution node is configured in NiFi.

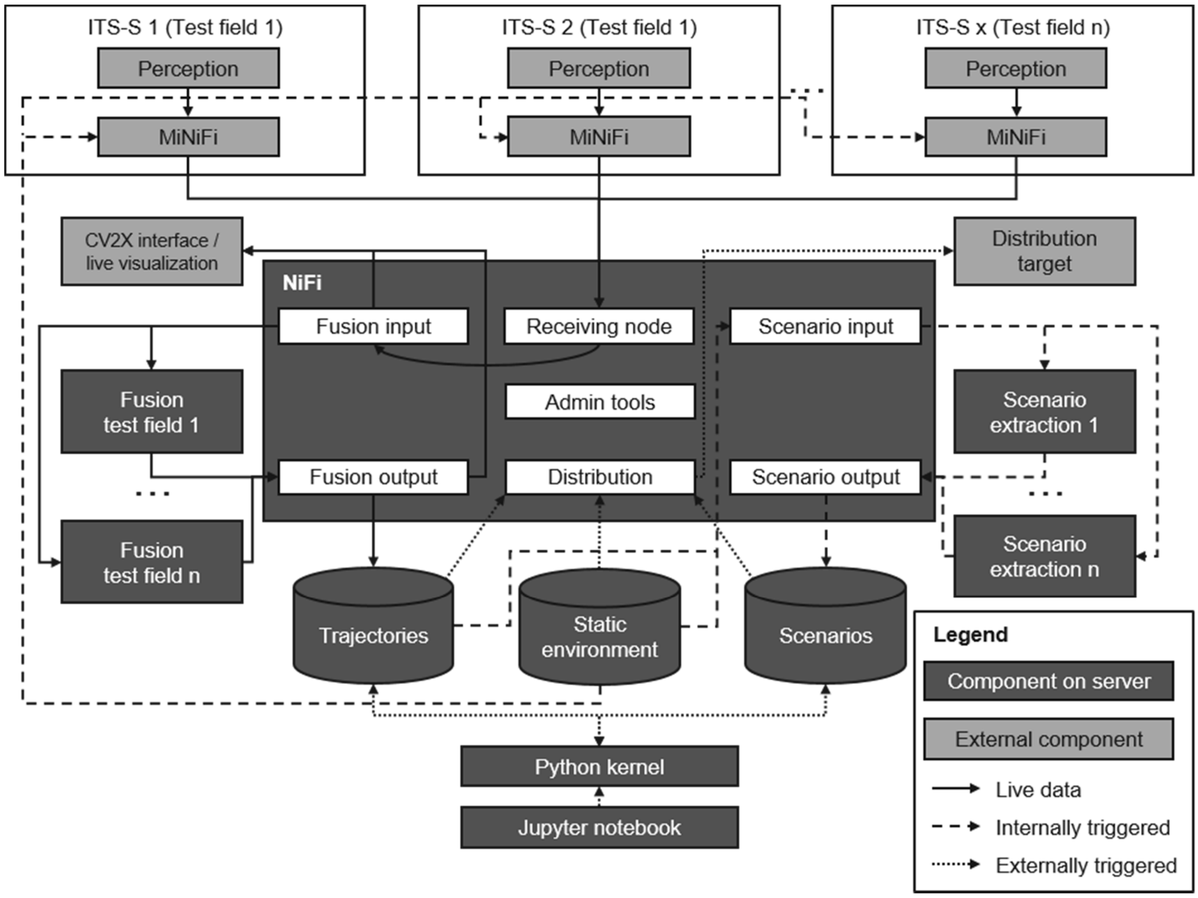

Fig. 13 Overview of data management concept.

4.4.2 Description of the Fusion Function

Fig. 14 shows the rough structure of the entire fusion function and in particular its division into several threads. A transmission control protocol (TCP) stream from NiFi serves as the input, in which all object lists of all measuring stations arrive and are collected. The first *receiver* thread has the task of decoding these object lists, converting them into the internal data format, and storing the object lists in a queue. The most important and time-critical thread is the *fuser*. It is merging the object lists and is storing all known objects in a global list.

Fig. 15 shows the process in more detail: In each iteration of the fuser, all object lists of the current queue are pulled and sorted by the timestamp of their included objects. This step is necessary because the stations' connections to the server operate on randomly varying delays and thus the incoming object lists don't necessarily arrive in the correct order. Thereupon, for each list, and thus for each time point of a measurement, all previously known objects are predicted to this time point using a predictive Kalman filter. The observed objects are associated with the known ones, their data is fused using the same filter and the updates are written back to the global object list. New objects that could not be associated with an existing one are created accordingly and provided with initial estimates for unmeasured quantities. In the case of the sensors installed so far, this concerns speed and acceleration; the former could also be measured directly, for example, by adding radar sensors to the measuring stations, which would also be taken into account accordingly in the filter. Each raw measured state is also passed to the Trajectory Updater, which writes it into the corresponding trajectory or creates a new one if necessary.

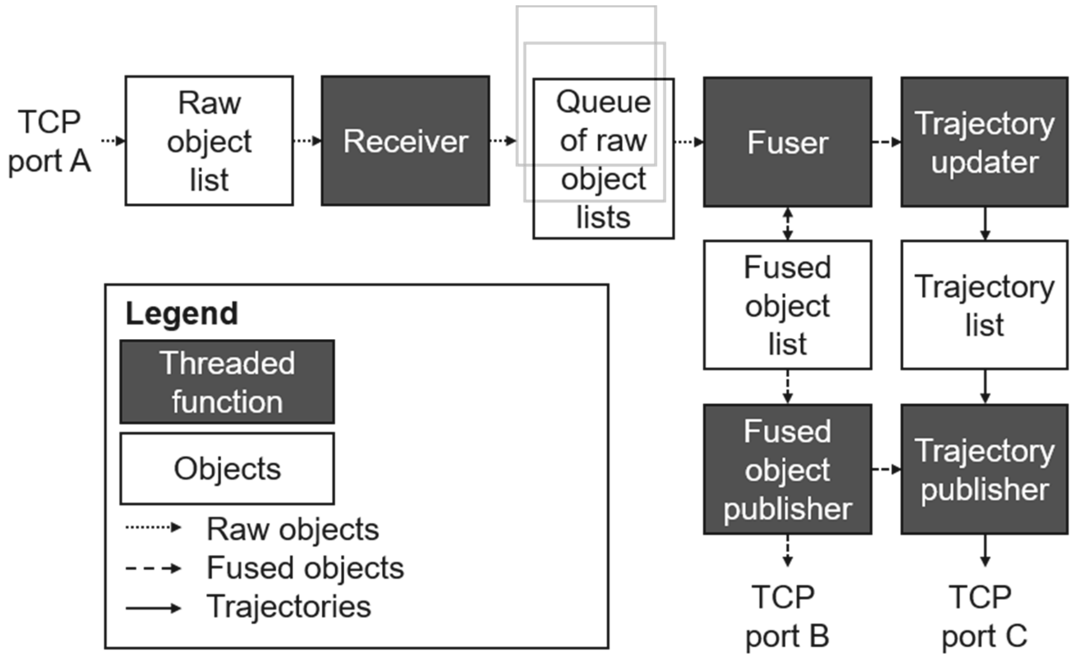

Fig. 14 Data processing pipeline of the data fusion function.

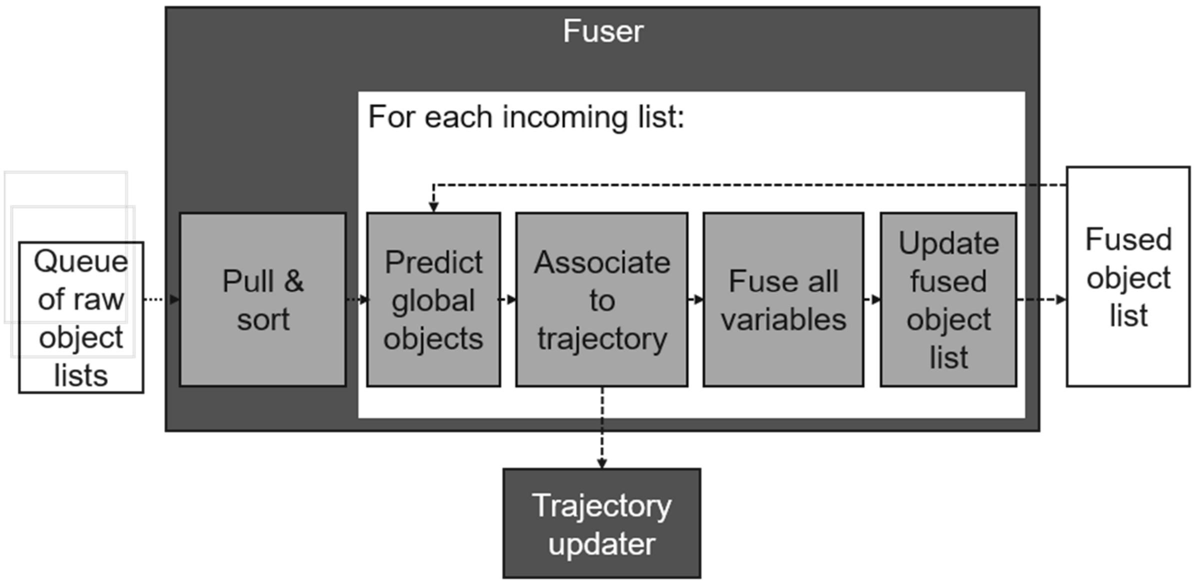

Fig. 15 Data processing pipeline of fuser function.

The right part Fig. 14 shows the generation of the output of the fusion function. It is presented in detail in Fig. 16 and consists of two threads: shown at the top is the Object Publisher, which predicts all currently known objects to the current timestamp at a fixed rate and publishes them back to NiFi via TCP. The data generated from this live fusion is also assigned to their corresponding trajectories. At the same time, the Object Publisher determines which objects have not been seen for longer than a certain threshold and signals the Trajectory Publisher accordingly. The latter takes the completed trajectory, applies an additional bidirectional filter to it for producing more accurate state estimates, and sends the trajectory to NiFi as well. Finally, all data associated with this object is deleted from the global lists. It should be noted that the trajectories must be collected in a queue during any connectivity problems so that the trajectory database is not missing any data in any case. For this reason, their publication runs in a separate thread, so that in case of problems no other processes are blocked. The situation is different for the object lists that are published. On the one hand, these serve as input for live visualization of the test fields and, on the other hand, as a source for a C-V2X stream, via which research vehicles in the test field learn what the infrastructure sensors have perceived. For both applications, sending object lists late due to internal problems may not only be useless but even potentially fatal if a CAV brakes because of an old object that has moved in the meantime. Thus, this TCP output is configured to discard the current object list in case of problems.

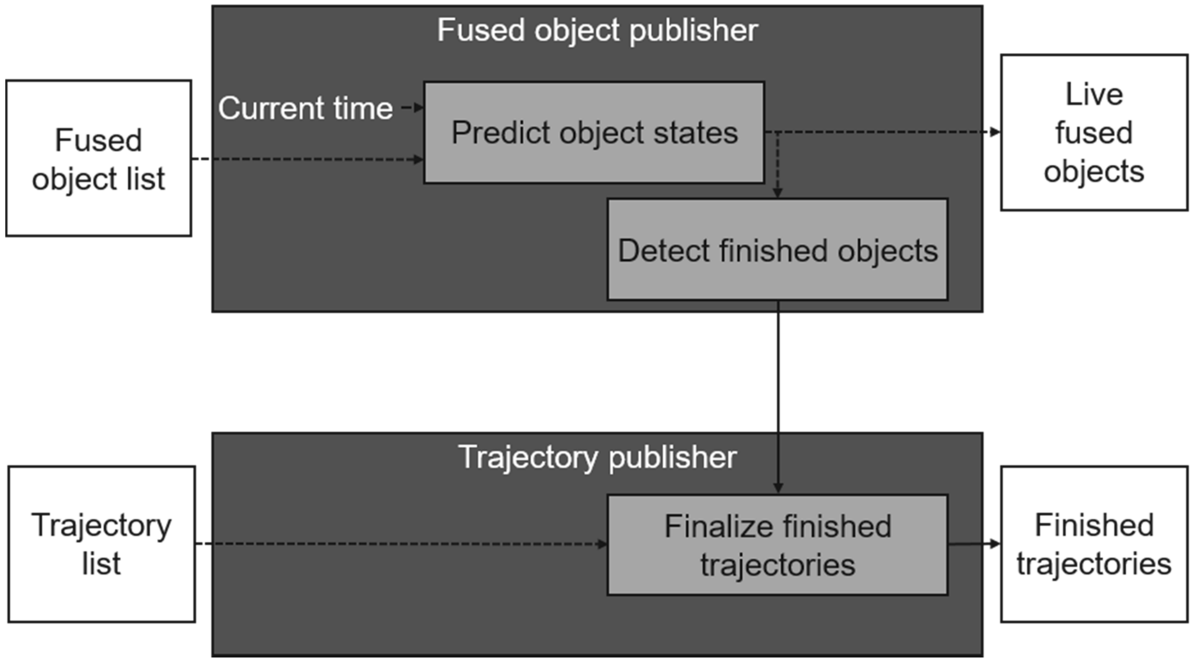

Fig. 16 Data processing pipeline for publishers of fused objects and trajectories.

In the case of computation time, the repetition rate of the thread *fuser* f_{Fuser} must be equal to or greater than the repetition rate of the publisher threads $f_{Publish} \leq 20$ Hz. Therefore, the greatest potential in terms of computation time optimization lies in the fuser thread.

4.5 Wireless Data Transfer

As described in the previous chapter, the collected object lists are provided at two different interfaces for V2X transmission. On the one hand, each ITS-S sends out its own locally collected object lists of each of the three sensor data streams and, on the other hand, the central computing server sends out a globally fused object list with the cumulated data of all ITS-Ss of one of the three test fields. Two different transmission methods are used. A distinction is made between data transmission via the automotive WLAN standard IEEE 802.11p [26], referred to as ITS-G5 by the European Telecommunications Standards Institute (ETSI) [27], and the cellular C-V2X approach, which uses 4G/5G technology. The main difference between ITS-G5 and C-V2X, in general, is the transmission range. ITS-G5 relies on the use of infrastructure-side RSUs and onboard units (OBUs). If an OBU is within the WLAN range of an RSU, it is able to receive its data. C-V2X, on the other hand, interacts within the known ranges of cellular networks.

With respect to the use of external infrastructure sensors for digital traffic recording, our work is limited to the standardized ETSI message type Collective Perception Message (CPM) [28]. The *collective perception* approach allows CAVs and/or intelligent infrastructure to exchange information about specifically perceived road users or objects with each other, which in particular extends the scope of the environment model of CAVs. In these cases, the object data to be exchanged is communicated in the form of CPMs.

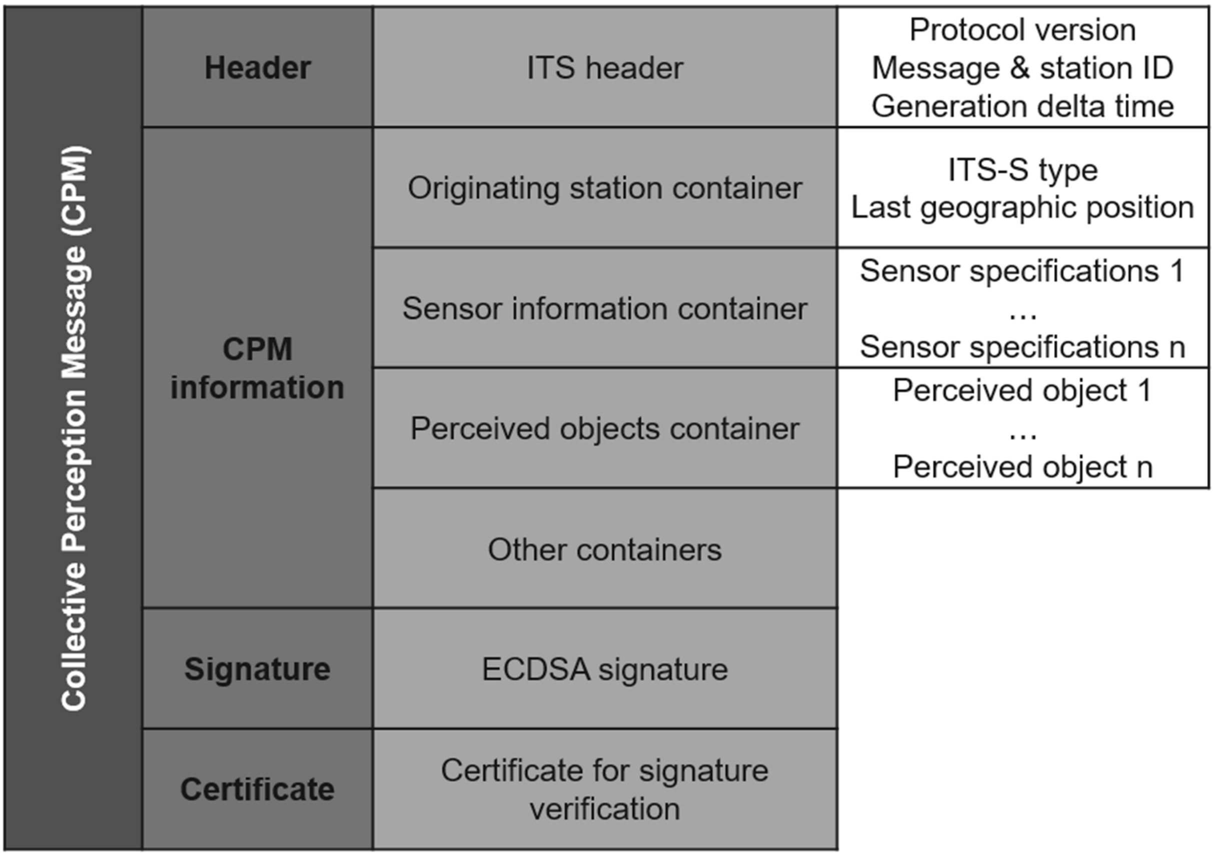

Fig. 17 General structure of a CPM [29].

Fig. 17 shows the general structure of such a CPM [29]. It consists of an intelligent transport systems (ITS) protocol data unit header and the actual CPM. The *ITS header* is the same for all standardized ETSI messages. The *originating station container* and the *sensor information container* contain information about the ITS-S type as well as about the installed sensor technology, such as its number, opening angle, and range. Finally, the *perceived objects container* provides all information about the objects detected by the sensors, such as classification, position, speed, heading, and dimensions. *Other containers* can in turn contain information about detected free areas and other less relevant information. Finally, in addition to the header and the containers, there is a *signature* in the form of an Elliptic Curve Digital Signature Algorithm (ECDSA) and a *certificate* to verify the signature.

4.6 Creation of a Digital Twin

All the work and development steps described above are the technical prerequisite for creating a digital twin. By combining intelligent connectivity and data collection, it is possible to transform a physical system into a digital twin. This contains the properties, states, and behavior of the real-world environment.

As described in section 4.4.2, the data fusion outputs a live object list at an approximately constant rate. This live data is the basis for an online digital twin of all test fields. Combined with the map data described in section 4.2, one can make sense of the perceived objects and for example, connect the data to a live visualization to observe their behavior. For this purpose, the accuracy of all data plays a less important role than the delay with which the live data is processed and presented, resulting in deviations of little more than 10 cm. A possible live visualization based on ROS and the Lanelet2 map is shown in Fig. 18. The data presented in the figure comes from simulations as no real-life data is available from the test fields yet. In the future, many more applications of the digital twin will become important. Notably, operators of CAVs are given a way to monitor the state of their fleet, interact with it, or even take over control in difficult situations.

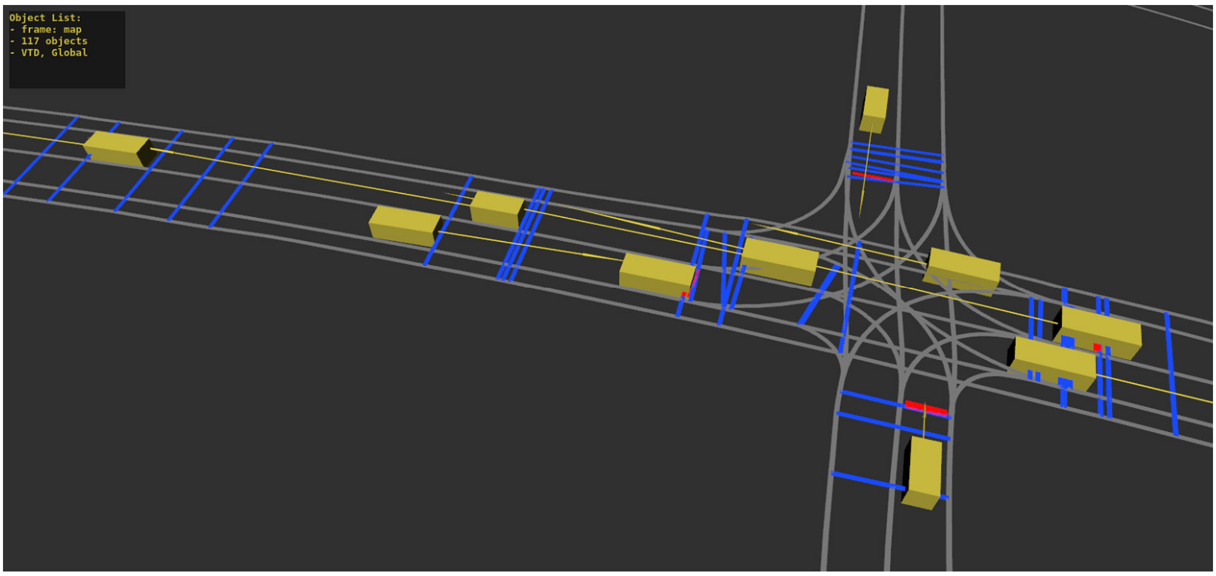

Fig. 18 A live visualization of the digital twin based on simulation data.

In addition, the digital twin can serve as a reference or extension for the onboard sensor system. Traffic data collected on the vehicle side can thus be compared or supplemented with the external environment data of the digital twin. This significantly increases the range of a CAV's virtual field of view and also avoids the problem of object occlusion. In particular, vulnerable road users (VRUs) that are hidden and not visible to the CAV, and whose safety should be the highest priority of any CAV, can thus be included in the current trajectory planning.

Besides utilizing the online digital twin, all vehicles at a certain period in time or a specific location can be reconstructed from the offline trajectory database. For this offline digital twin, the trajectory of each vehicle is once again filtered using bidirectional algorithms. These lead to improved accuracy of usually less than 10 cm standard deviation. The offline data can mainly be used for the testing of automated driving functions in a realistic environment or for the parameterization of driver models for simulations.

5 Conclusion and Outlook

This paper shows all steps of the value chain within the ACCorD research project to generate a highly accurate digital twin of three digital test field sections in urban, rural, and highway environments. The digital test fields are capable of detecting all surrounding vehicles and VRUs. Raw data from infrastructure sensors at the roadside of the test fields are processed into object data using complex algorithms and deep learning approaches. These are in turn forwarded in real-time to a central computing server, which processes the data and forwards it both to a trajectory database and to research vehicles on the corridor. On the vehicle side, for example, the measurement data can thus act as a reference for evaluating the internally collected environmental data. The constantly filling trajectory database can in turn be used for further research and development purposes, such as the extraction and analysis of user-defined scenarios.

Although planning and setting up such digital test fields is a time-consuming and cost-intensive challenge, it subsequently offers the valuable possibility of permanent traffic data collection. The collected traffic data can be imported into a digital twin both online and offline. In particular, by importing highly accurate measurement data with an achieved standard deviation of about 10 centimeters, the digital twin offers the possibility to develop and safeguard automated and connected driving functions in real traffic in the best possible way. With the completion and opening of the test fields in the course of this year, a unique test environment will be created in North Rhine-Westphalia, paving the way for numerous innovative research and development activities in the coming years.

In the future, tests will also be conducted to determine to what extent the use of 5G will lead to a further reduction in data transmission latencies between intelligent infrastructure sensors and automated and connected vehicles compared with 4G and ITS-G5. Only with the lowest possible latencies can safety-relevant traffic data be used profitably in an online digital twin. In addition, research will be conducted to determine the extent to which the general energy consumption of intelligent infrastructure sensors can be reduced and optimized by using field-programmable gate arrays (FPGAs) to evaluate sensor raw data.

6 Acknowledgment

The research leading to these results is funded by the Federal Ministry of Transport and Digital Infrastructure (BMVI) within the project "ACCorD: Corridor for New Mobility Aachen - Düsseldorf" (FKZ 01MM19001A). The authors would like to thank the consortium for the successful cooperation.

7 Abbreviations

ACCorD	Corridor for new Mobility Aachen - Düsseldorf
AIM	Application Platform Intelligent Mobility
ATC	Aldenhoven Testing Center
ASAM	Association for Standardization of Automation and Measuring Systems
BMVI	Federal Ministry of Transport and Digital Infrastructure
CAD	Computer-Aided Design
CAV	Connected and Automated Vehicle
CPM	Collective Perception Message
CPU	Central Processing Unit
C-V2X	Cellular-V2X
DLR	German Aerospace Center
ECC	Error Correction Code
ECDSA	Elliptic Curve Digital Signature Algorithm
ESRI	Environmental Systems Research Institute
ETSI	European Telecommunications Standards Institute
FOV	Field Of View
FPGA	Field-Programmable Gate Array
GNSS	Global Navigation Satellite System
GPU	Graphics Processing Unit

IMU	Inertial Measurement Unit
ITS	Intelligent Transport Systems
ITS-S	Intelligent Transport Systems Station
OBU	Onboard Unit
PoE	Power-over-Ethernet
PTRZ	Pan, Tilt, Roll, Zoom
RAID	Redundant Array of Independent Disks
RAM	Random-Access Memory
ROS	Robot Operating System
RSU	Roadside Unit
SLAM	Simultaneous Localization And Mapping
TCP	Transmission Control Protocol
V2X	Vehicle-to-everything Communication
VRU	Vulnerable Road User

8 References

[1] Rose, Michael, Wanner, Matthias, Hilger, Annaliesa, 2018.
Das Reallabor als Forschungsprozess und -infrastruktur für nachhaltige Entwicklung. Konzepte, Herausforderungen und Empfehlungen.
In: Nachhaltiges Wirtschaften - NaWiKo Synthese Working Paper No. 1.
Karlsruhe: Fraunhofer Institute for Systems and Innovation Research ISI, pages 1-36.

[2] Federal Ministry of Transport and Digital Infrastructure (BMVI), 2021.
Übersicht zu Digitalen Testfeldern für das automatisierte und vernetzte Fahren im Realverkehr in Deutschland.
Online [Access 13-06-21]:
https://www.bmvi.de/SharedDocs/DE/Anlage/DG/Digitales/uebersicht-digitale-testfelder-avf-bmvi.pdf?__blob=publicationFile

[3] Wang, Jian, Shao, Yameng, Ge, Yuming, Yu, Rundong, 2019.
A Survey of Vehicle to Everything (V2X) Testing.
In: Sensors 2019. 19, 334, pages 1-20.
doi: 10.3390/s19020334

[4] de Gelder, Erwin, Paardekooper, Jan-Pieter, 2017.
Assessment of Automated Driving Systems Using Real-Life Scenarios.
In: 2017 IEEE Intelligent Vehicles Symposium (IV). Redondo Beach (CA), USA, June 11-14, 2017.
New York, NY, USA: IEEE, pages 589-594.
doi: 10.1109/IVS.2017.7995782

[5] ACCorD Consortium, 2021.
Corridor for New Mobility Aachen – Düsseldorf (ACCorD).
Website: https://www.accord-testfeld.de/

[6] Kloeker, Laurent, Kloeker, Amarin, Thomsen, Fabian, Erraji, Armin, Eckstein, Lutz, Lamberty, Serge, Fazekas, Adrian, Kalló, Eszter, Oeser, Markus, Fléchon, Charlotte, Lohmiller, Jochen, Pfeiffer, Pascal, Sommer, Martin, Winter, Helen, 2021.
Corridor for new mobility Aachen - Düsseldorf: Methods and concepts of the research project ACCorD.
In: arXiv preprint.

[7] KoMoDnext Consortium, 2021.
KoMoDnext – Automated Driving in the Digital Test Field Düsseldorf.
Website: https://komodnext.org/

[8] CERMcity Consortium, 2021.
Center for European Research on Mobility Urban Validation Environment (CERMcity).
Website: https://www.ika.rwth-aachen.de/de/forschung/projekte/automatisiertes-fahren/2718-cermcity.html

[9] Providentia++ Consortium, 2021.
Proactive Video-Based Use of Telecommunication Technologies in Innovative Traffic Scenarios (Providentia).
Website: https://innovation-mobility.com/

[10] Kraemmer, Annkathrin, Schoeller, Christoph, Gulati, Dhiraj, Knoll, Alois, 2019.
Providentia - A Large Scale Sensing System for the Assistance of Autonomous Vehicles.
In: Robotics: Science and Systems (RSS), Workshop on Scene and Situation Understanding for Autonomous Driving. Freiburg, June 22, 2019.

[11] DIGINET-PS Consortium, 2021.
Digitally Connected Protocol Route (DIGINET-PS).
Website: https://diginet-ps.de/

[12] Rakow, Christian, Khan, Manzoor Ahmed, 2018.
Mobility as a Service Enabled by the Autonomous Driving.
In: International Conference on Internet of Vehicles. Paris, France, November 20-22, 2018.
Cham, Switzerland: Springer International Publishing, pages 208-219.
doi: 10.1007/978-3-030-05081-8_15

[13] German Aerospace Center (DLR), 2021.
Test Bed Lower Saxony for automated and connected mobility.
Website: https://verkehrsforschung.dlr.de/de/projekte/testfeld-niedersachsen-fuer-automatisierte-und-vernetzte-mobilitaet

[14] Köster, Frank, Mazzega, Jens, Knake-Langhorst, Sascha, 2018.
Automatisierte und vernetzte Systeme Effizient erprobt und evaluiert.
In: ATZextra. 23(5), pages 26-29.
doi: 10.1007/s35778-018-0040-9

[15] Institute for Automotive Engineering (ika), RWTH Aachen University, 2021.
High-precision digital traffic recording as a basis for future mobility research - Construction of mobile and modular measuring stations (HDV-Mess).
Website: https://www.ika.rwth-aachen.de/en/research/projects/automated-driving/2929-hdv-mess.html

[16] Kloeker, Laurent, Kloeker, Amarin, Thomsen, Fabian, Erraji, Armin, Eckstein, Lutz, 2020.
Traffic Detection Using Modular Infrastructure Sensors as a Data Basis for Highly Automated and Connected Driving.
In: 29th Aachen Colloquium Sustainable Mobility 2020. Aachen, October 5-7, 2020.
Aachen: Aachener Kolloquium Fahrzeug- und Motorentechnik GbR, pages 1835-1844.
ISBN: 978-3-00-064871-7

[17] Kloeker, Laurent, Thomsen, Fabian, Eckstein, Lutz, Trettner, Philip, Elsner, Tim, Nehring-Wirxel, Julius, Schuster, Kersten, Kobbelt, Leif, Hoesch, Michael, 2021.
Highly accurate digital traffic recording as a basis for future mobility research: Methods and concepts of the research project HDV-Mess.
In: arXiv preprint arXiv:2106.04175.

[18] German Aerospace Center (DLR), 2021.
Application Platform for Intelligent Mobility (AIM).
Website: https://www.dlr.de/ts/en/desktopdefault.aspx/tabid-6422/#gallery/25304

[19] Google LLC, 2021.
Google Maps.
Website: https://www.google.de/maps

[20] ESRI Inc., 2021.
Shapefile.
Website: https://www.esri.com/content/dam/esrisites/sitecore-archive/Files/Pdfs/library/whitepapers/pdfs/shapefile.pdf

[21] ASAM e.V., 2021.
ASAM OpenDRIVE.
Website: https://www.asam.net/standards/detail/opendrive/

[22] IPG Automotive GmbH, 2021.
Import of road descriptions into CarMaker.
Website: https://ipg-automotive.com/de/support/kundenbereich/faq/ticket/import-of-road-descriptions-into-carmaker/

[23] Poggenhans, Fabian, Pauls, Jan-Hendrik, Janosovits, Johannes, Orf, Stefan, Naumann, Maximilian, Kuhnt, Florian, Mayr, Matthias, 2018.
Lanelet2: A High-Definition Map Framework for the Future of Automated Driving.
In: 2018 IEEE Intelligent Transportation Systems Conference (ITSC). Maui (HI), USA, November 4-7, 2018.
New York, NY, USA: IEEE, pages 1672-1679.
doi: 10.1109/ITSC.2018.8569929

[24] Ouster, Inc., 2021.
Ouster OS1 product website.
Website: https://ouster.com/products/os1-lidar-sensor/

[25] Bosch Sicherheitssysteme GmbH, 2021.
Bosch Flexidome IP Starlight 8000i User Manual.
Website: https://resources-boschsecurity-cdn.azureedge.net/public/documents/FLEXIDOME_IP_starlig_Data_sheet_deDE_68669614475.pdf

[26] IEEE - Institute of Electrical and Electronics Engineers, 2010.
802.11p-2020 - IEEE Standard for Information Technology-- Telecommunications and Information Exchange Between Systems-- Local and Metropolitan Area Networks-- Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications; Amendment 6: Wireless Access in Vehicular Environments.
New York, NY, USA: IEEE

[27] Festag, Andreas, 2014.
Cooperative Intelligent Transport Systems Standards in Europe.
In: IEEE Communications Magazine. Vol. 52, No. 12, pages 166-172.
doi: 10.1109/MCOM.2014.6979970

[28] ETSI - European Telecommunications Standards Institute, 2019.
ETSI TR 103 562 V2.1.1 - Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Analysis of the Collective Perception Service (CPS); Release 2.
Sophia Antipolis Cedex, France: ETSI

[29] Schiegg, Florian A., Bischoff, Daniel, Krost, Johannes R., Llatser, Ignacio, 2020.
Analytical Performance Evaluation of the Collective Perception Service in IEEE 802.11p Networks.
In: 2020 IEEE Wireless Communications and Networking Conference (WCNC).
Seoul, Korea (South), May 25-28, 2020.
New York, NY, USA: IEEE, pages 1-6.
doi: 10.1109/WCNC45663.2020.9120490