N. Druml et al., "Programmable Systems for Intelligence in Automobiles (PRYSTINE): Final results after Year 3,"
2021 24th Euromicro Conference on Digital System Design (DSD), 2021, pp. 268-277, doi: 10.1109/DSD53832.2021.00049.

Programmable Systems for Intelligence in Automobiles (PRYSTINE): Final results after Year 3

Norbert Druml, Anna Ryabokon, Rupert Schorn, Jochen Koszescha, Kaspars Ozols, Aleksandrs Levinskis, Rihards Novickis, Ethiopia Nigussie, Jouni Isoaho, Selim Solmaz, Georg Stettinger, Sergio Diaz, Mauricio Marcano, Jorge Villagra, Juan Medina, Martina Schwarz, Antonio Artuñedo, Mauro Comi, Rutger Beekelaar, Onur Özçelik, Elif Aksu Taşdelen, Yeşim Gürbüz, Jan Saijets, Jukka Kyynäräinen, Dmitry Morits, Björn Debaillie, Maxim Rykunov, Joan Escamilla, Jarno Vanne, Tomi Korhonen, Kalle Holma, Eva-Maria Matzhold, Carlo Novara, Fabio Tango, Paolo Burgio, Giuseppe Calafiore, Milad Karimshoushtari, Emilie Boulay, Miguel Dhaens, Kylian Praet, Han Zwijnenberg, Henri Palm, David Aledo Ortega, Ercan Kalali, Tuomas Pensala, Arto Kyytinen, Morten Larsen, Omar Veledar, Georg Macher, Michael Lafer, Lorenzo Giraudi, Jakob Reckenzaun, Daniel Hammer, Naveen Mohan, Josef Schmid, Alfred Höß, Shai Ophir, Anand Dubey, Jonas Fuchs, Maximilian Lübke, Andrei Anghel, Nicolae-Cătălin Ristea, Martin Törngren, Alua Musralina, Marlene Harter, Joseena Memadathil Jose, George Dimitrakopoulos

***Abstract*— Autonomous driving is disrupting the automotive industry as we know it today. For this, fail-operational behavior is essential in the sense, plan, and act stages of the automation chain in order to handle safety-critical situations on its own, which currently is not reached with state-of-the-art approaches.**

The European ECSEL research project PRYSTINE realizes Fail-operational Urban Surround perceptION (FUSION) based on robust Radar and LiDAR sensor fusion and control functions in order to enable safe automated driving in urban and rural environments. This paper showcases some of the key exploitable results (e.g., novel Radar sensors, innovative embedded control and E/E architectures, pioneering sensor fusion approaches, AI-controlled vehicle demonstrators) achieved until its final year 3.

***Keywords*— FUSION, fail-operational, perception**

I. Introduction

PRYSTINE (PRogrammable sYSTems for INtelligence in AutomobilEs) realizes Fail-operational Urban Surround perceptION (FUSION) by researching and developing a set of key technologies. On lowest abstraction layer, novel components are researched and developed, such as robust Radar and LiDAR sensors, ASIL-D safety controllers, and number-crunching hardware accelerators for CNNs. Using these components, PRYSTINE realizes the next generation autonomous driving platforms providing both ASIL-D embedded control and high-performance processing power. Dependable embedded control and embedded intelligence is then achieved by co-integration of signal processing and AI approaches. These novel sensor fusion and decision making solutions build the foundation for PRYSTINE's vehicular demonstrators.

This work presents in detail the latest research achievements and is structured as follows:

- Section II and III present innovative results on component level, in particular in the fields of LiDAR sensors, Radar sensors, RF interference mitigation, and hardware accelerators for CNNs.
- Section IV and V showcase PRYSTINE's results in the field of fail-operational E/E architectures and central embedded intelligence platforms featuring ASIL-D and providing the computing performance for FUSION and applications.
- Section VI, VII, and VIII present groundbreaking achievements in the field of sensor fusion, decision making, and vehicular demonstrators (such as passenger cars, heavy duty trucks semi- and full-sized trailers), and AI controlled co-pilots.

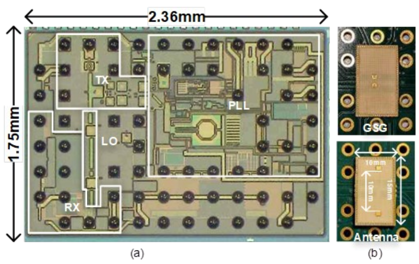

Fig. 1: The 60GHz scalable radar (a) chip micrograph and (b) its module with antennas and GSG pads.

II. Next Generation Radar and LiDAR Sensors

Automotive Radar: automated driving (AD) relies on many sensors to perceive the environment around the vehicle. In general, these sensors are designed for maximal performance to obtain the best possible view. However, this continuously stresses the computing and battery resources. This strategy might be reviewed in critical situations when, for example, the central battery is failing and the remaining power is focused to drive the vehicle to a safe harbor. In such scenario, electrical systems should reconfigure to consume minimal power while offering basic performance. In PRYSTINE, the consortium developed a 60 GHz radar sensor (see Fig. 1) that supports such fail-operationality: the sensor offers a high degree of power scalability by duty-cycling its operation. The sensor's ability to "instantaneously" de-/re-activate allows to insert "micro-sleeps" and to reduce its average power consumption. Validation experiments of the developed radar sensor show an acceptable sensitivity (10dB SNR) up to 6.25% duty cycling, which corresponds to a factor of 10 in power reduction. The reduced duty-cycling however reduces the radar's maximum unambiguous velocity, increases the capturing and processing time, and makes the system more vulnerable to interferences. The power versus performance can be traded off depending on the fail incident criticality as the duty-cycling factor is continuously scalable.

Apart from novel ASIC developments, based on the PRYSTINE's year 2 results [1], the consortium designed a novel Fully Convolutional Network (FCN) for Radio Frequency Interference (RFI) mitigation that can recover the phase along with the magnitude of radar beat signals and can cope with multiple non-coherent RFI sources. The proposed network takes as input the real part, the imaginary part and the magnitude of the Short-Time

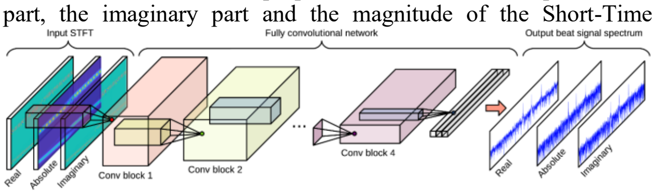

Fig. 2: Architecture of the FCN model. The input STFT is processed through a series of four conv blocks (composed of conv and pooling layers) until the vertical dimension is reduced to 1, while preserving the horizontal dimension. The output is the complex range profile without interference.

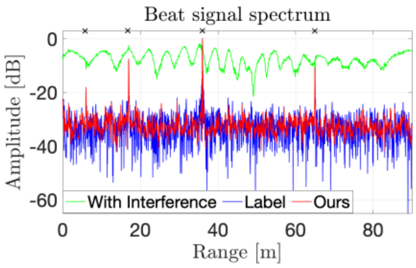

Fig. 3: Sample range profile (magnitude) for a case with 2 interference sources.

Fourier Transform (STFT) of the beat signal with interference, providing as output the real part, the imaginary part and the magnitude of the range profile, respectively. The developed FCN model was compared with some baseline methods in a series of comprehensive experiments on simulated and real data, showing that the proposed approach provides superior results, especially when multiple interference sources are present. The general architecture of the network and a sample result for 2 interference sources is shown in Fig. 2 and Fig. 3.

In contrast to state-of-the-art having independent classification and tracking, PRYSTINE's consortium developed a Bayesian integrated framework for combined detection, classification and tracking of vulnerable road users using radar. The tracker's state vector is modified to combine both localization parameters and feature embedding vector, as the temporal model and appearance model of targets respectively. The feature embedding is learned by an auto-encoder based CNN architecture, as shown in Fig. 4. As consequence, the tracker's performance is optimized due to better separability of the targets and further enhanced due to Bayesian formulation utilizing the temporal smoothing of the classifier's embedding vector.

Automotive LiDAR: Advanced Driver Assistance Systems (ADAS) and self-driving vehicles use laser range finding systems for mapping their surroundings and providing data for decision-making. Pressing down the size and the cost of the system, while maintaining the high performance (long range, high angular resolution, and high scan rate) presents a significant challenge.

A key element of a LiDAR is the scanning system. In PRYSTINE, the partners developed high-amplitude omnidirectional 2D piezo MEMS scanning mirrors. Resonant operation with vacuum encapsulation enables up to ±16° mechanical tilt angles with low drive voltages of 1V to 4V, depending on the design. The power consumption of the micromirrors with the control electronics is lower than 100 mW. The mirror drive electronics uses square wave for driving and pulse width modulation to adjust the power instead of amplitude. Delta connection increases the voltage over actuator compared to driving each actuator individually. As a significant advantage, piezoelectric sensing provides feedback signals linearly proportional to the tilt angles of the micromirrors. The mirror technology is applicable to a wide range of LiDAR applications besides ADAS.

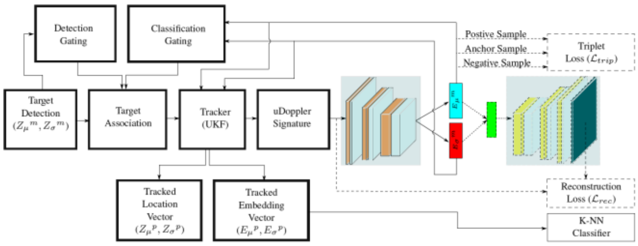

Fig. 5: Illustration of the training (dotted lines) and inference (solid lines) phase of the developed integrated Bayesian radar classification and tracking framework [2].

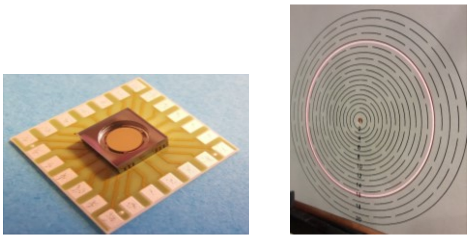

Fig. 4: Vacuum encapsulated 4 mm diameter mirror (left) driven with an omnidirectional scanning pattern with 15.7° mechanical tilt angle (right).

III. Hardware Accelerators for CNN

Artificial Intelligence (AI) is key for automated vehicles, in particular Deep Neural Networks (DNN), which are composed of a different number of convolutional layers. These Convolutional Neural Networks (CNNs) are used for signal processing, data fusion, object detection and classification, image segmentation, and other operations vital for AD. Current implementations rely on Graphic Processor Units (GPUs), which provide the computation throughput needed for real-time applications. However, GPUs consume a great amount of power, reducing considerably the autonomy of the vehicle. One of PRYSTINE's objectives is to provide a hardware accelerator capable of producing the same amount of computation throughput, but with reduced power consumption.

Among all the possible accelerator architectures, a systolic array architecture has been selected. It provides a reasonable trade-off between computation parallelism, data re-use, required memory bandwidth, and other critical design parameters. Systolic arrays for CNN acceleration are not new. Some good examples are Google's TPU [3] or Eyeriss [4]. However, most GPUs have been designed (and optimized) for usage in data-centers, which require the GPUs to work as clients on a host processor. PRYSTINE's systolic array has been designed to accept several data-streams from different sensors and to perform data fusion. It processes the data in a pipeline/data-flow fashion automatically, without the intervention of a host unless required. Additionally, it provides a kind of plug-and-play Processing Elements (PE) that allow computation on different possible data types for inputs and weights, approximate computations, quantization, etc. The array is divided into 3 x 3 systolic array sub-tiles, from which one column of them (3 x 1) can be of a different kind, allowing hybrid solutions for different sensor inputs and/or different precision for hidden layers (the same accuracy can be obtained reducing the precision of the hidden layers if the first and last layers keep their original precision).

Although it has been designed for an ASIC realization, the first prototype was implemented in an FPGA (Xilinx's Zynq SoC). Fig. 6 shows the prototype's system architecture. It receives data from 3 different sensors: camera, radar, and LiDAR. Sensor data is synchronized and streamed via AXI4 streams that feed the IMem buffer(s) of the systolic array. Weights will be loaded into the WMem buffer from the off-chip DRAM. The outputs are streamed from the OMem buffer.

The systolic array is integrated into the TVM software stack [5]. This will allow DNNs to be implemented in popular libraries like TensorFlow, PyTorch, Keras, etc., to be compiled for inference into

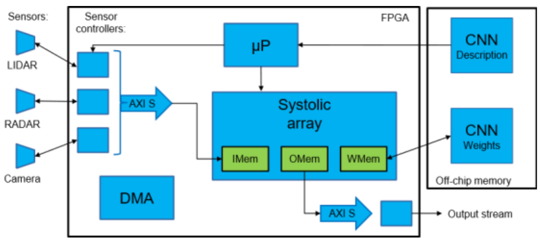

Fig. 6: System architecture of the CNN hardware accelerator.

the accelerator.

IV. High Performance Embedded Control and Intelligence for FUSION

In the presence of the ever-increasing connectivity inside and outside modern vehicles as well as of the large amount of sensor data to be processed, the complexity of E/E systems and their software functionality is growing exponentially. Safe and secure communication required for ADAS and AD functions is one of the most critical challenges in realizing highly automated driving. One of the main objectives of PRYSTINE is to push the technological capabilities for AD upwards from SAE Level 2 "Partial Driving Automation" for on-road vehicles to SAE Level 3+ "Conditional/High Automation", by *implementing fail-operational mechanisms for embedded control and enhancing sensor fusion by 360-degree perception using cameras, LiDAR, and Radar*. This objective is considered on a system-level by addressing a set of associated KPIs, each implemented and recently validated in the specific solutions/demonstrators described below. PRYSTINE's results are being successfully integrated in three types of vehicles: (i) passenger vehicles *KIA SOUL EV* [6] *and Ford Mondeo*, (ii) *Ford heavy duty vehicle* and (iii) *Renault heavy-duty truck with trailer*, representing an excellent basis for the upcoming real-world validation of several complex requirements and exploitation of the project achievements.

Fail-operational autonomous driving platform enhances existing architectures, addresses the huge complexity of the E/E systems and implements AD functions for future automated mobility while ensuring the highest automotive safety standards (ISO26262 up to ASIL D) and state-of-the-art high-speed connectivity. This novel architecture (see Fig. 7) benefits from a modularity concept combining COTS elements such as SoCs, AURIX™ automotive microcontroller, power supply, Deterministic backbone network for low-latency data exchange, multiple cameras, etc. and allows for flexibility and market speed-up of the developed solution. The implemented fail-over mechanism realizes embedded control in a non-compromised way to advance safety. The project's consortium is closely cooperating in this project to bring the ADAS architectures to next level of safety enabling highly automated driving.

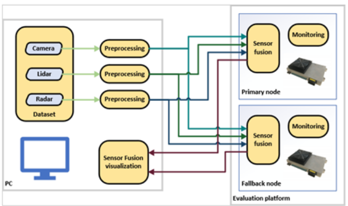

Fig. 7: Controller redundancy.

AI-based perception system able to perform data fusion, processing and decision making is based on the developed COMPAGE framework (component-based software management framework) and DL-based algorithms capable of identifying faulty sensors by analyzing data of different types, i.e., LiDAR, Radar, cameras. PRYSTINE's partners developed their version of the OSCC based DriveByWire (DbW) system (see Fig. 8) and ICOM communication framework which enables efficient inter-component data transfer (zero-copy, i.e., memory contents are not duplicated) and synchronization, making it a good candidate for a fault-tolerant DDS alternative, and thus, valuably contributes to the development of distributed platforms. Fail-operational object tracking and fusion have been tested in real scenarios, involving RADAR, LiDAR and camera sensors as well as the DbW system along with the other components. The following DL methods are incorporated in the proposed solution: (a) Point cloud object recognition utilizing DNNs, i.e. PointPillar, (b) LiDAR-based localization and path execution for operation in an urbanized environment, and (c) camera-based inverse perspective mapping for object tracking and registration in a coherent bird's eye view [7].

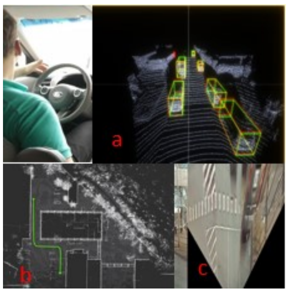

Fig. 8: DwC demonstrator

Fig. 9: Perception system prototype.

A perception system with reliable and secure data was developed to guarantee accurate perception despite faults and security threats in the sensors and computing platform. Unlike existing solutions that focus on tackling security issues individually, the developed approach (see Fig. 9) is based on a fully integrated security engineering process encompassing a customized firewall, intrusion detection system and trust model for evaluating the trustworthiness of sensors data. In addition to enhanced security, the integration of the trust model with sensor fusion improves the accuracy of object detection and tracking. To assist automated heavy-duty vehicle driving, fail-operational middleware that detects faults and security incidents in both software and hardware platforms of the controller was developed. The overall simulation of the perception system is being carried out using real data collected by heavy-duty truck in different weather conditions. This is necessary to improve the accuracy of the trust model, the middleware, intrusion detection system, and ultimately for realizing secure data communication and sensors' reliability-aware data fusion. The goal is to achieve 95% successful security threats and faults detection in sensors and computing platform.

Autonomous parking solution deploying prototype FUSION algorithms for low-speed autonomy is a novel approach that contributes in the field of Automated Parking Vale (AVP) Systems both for passenger and commercial vehicle segments. Fusion algorithms and perception components were utilized to achieve an SAE Level 3+ equivalent autonomous parking and low-speed autonomy solution, which provides fail-operationality and robustness by the utilization of multiple sensor sources (e.g. cameras, LiDAR, and Radar) together with novel sensor fusion algorithms. The developed FUSION components and software blocks were ported to the Autoware.AI software stack for validation on the Automated Driving Demonstrator (ADD) passenger vehicle platform (Ford Mondeo) using a set of validation scenarios (e.g. back maneuver into a selected parking slot within a possibly dynamic environment). The tests were successfully completed at the ÖAMTC Lang/Lebring proving ground in Graz, Austria. In the final phase of the project, relevant perception components are being ported and integrated into the Ford heavy-duty truck for comparable automated back-parking use-case demonstrations.

V. Fail-Operational E/E Architecture Enabling FUSION

PRYSTINE looks into suitable vehicle architectures and communications to accommodate the need for automated vehicles to sense their environment and respond accordingly [8]. Three demonstratable aspects are covered, as described in this section. Their integration phase is based on the reported components [1].

The first fail-operational E/E architecture demonstrator is a direct response to the increasing need for integration of centralized adaptive control in automotive systems [9]. The architecture facilitates communication between the optimized fail-operational sensor combinations and dependable embedded controllers. The fail-operational behavior in this scope is ensured using redundancy at the hardware level, which forms the basis for the integration of optimized vehicular E/E infrastructure and communication systems [9]. The hardware arrangement enables periodic status messages to determine the health of the duplicated electronic controllers. The messaging utilizes TSN Ethernet features, in particular, 802.1CB aka "Frame Replication and Elimination for Reliability", which leverages two communication paths. A failure of a single communication channel or the embedded processor does not limit the operation of the system, as the failing control branch is replaced by the still functioning and connected processor. In the event of the entire control unit failure, the redundant controller takes over the control. In such a case, some sensor inputs may become unusable, subject to the active configuration (as depicted in Fig. 10).

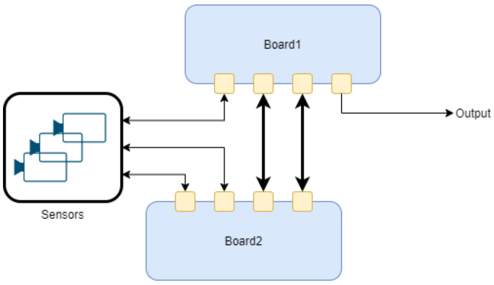

Fig. 10: Controller redundancy.

When considering communication to the external devices (vehicles, infrastructure etc.), the performance of Bluetooth Low Energy 5 (BLE5) advertisement modes is analyzed with other communication technologies, to define recommendations for integration of the specific gateway protocols.

At the sensor level, duplication approach is also the most promising answer. This includes duplication of the object detection/tracking pipelines of the same type as well as the duplication of different object detection/tracking algorithms, which rely on the sensor input data. A representative visualization snapshot of the object tracking algorithms in real-world driving is depicted in Fig. 11.

The second fail-operational E/E architecture demonstration considers fault tolerant functional safety in real-world driving scenarios using a testbed, which is resulting from the PRYSTINE project. The AD-EYE [10] testbed provides simulation, physical and mixed reality modes of operation, as well as test automation. It enables nominal, safety and security testing scenarios (e.g. fault injection and cyber-attacks). The provided benefits rely on mixed-mode setups, from simulation over the vehicle in the loop to fully physical setups. In this particular demonstration, the testbed exploits system modelling as a Semi-Markov process to characterize and map the influence of environmental scenarios and states on the system-level behaviour. A direct result of such an approach is an improved instantiation of the supervisory architecture pattern (Fig. 12), which is designed to maximize the availability of the system without violating the constraints set by the safety analysis. Additional results are based on open-source data and tools. The first instance of open-source tools (the SMP-TOOL) and data (the OPEN-KTH dataset) are already released in practice. The ambition is to release the entire testbed, associated design, components and code as open-source.

The third E/E demonstration investigates mobile network communications, which are exploited by AD for safe and reliable transmission of large amounts of data, such as updates of high-precision digital maps. Since the accurate detection of the environment and a current high-precision digital map are basic requirements for autonomous navigation, the mobile network connection must provide the highest possible data rates with the

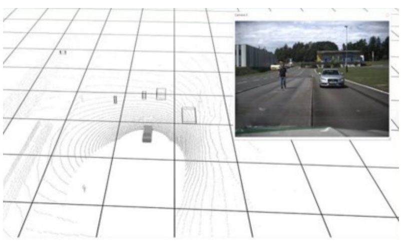

Fig. 12: Object tracking visualisation under driving test-track conditions.

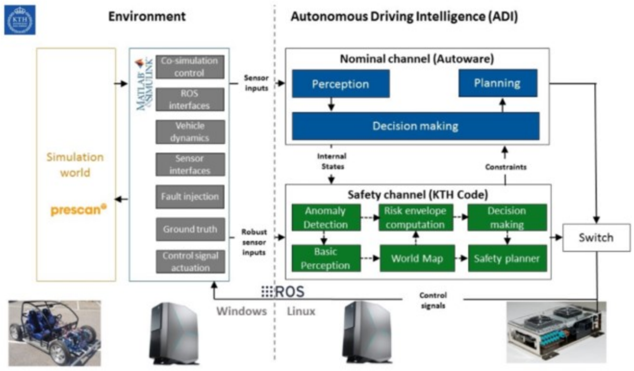

Fig. 11: AD-Eye: a testbed for Automated Driving and ITS.

lowest possible latency times. As mobile phone coverage is not always guaranteed, especially in rural areas and between national borders, the arising issues are signal coverage gaps and longer switching times between mobile phone towers. This has the potential to destabilize connections and cause high latency during autonomous driving. To avoid this realistic risk and to increase the connection quality of the communication between the vehicle and a server, AI-based methods are developed to predict the future latency of a mobile network connection. The demonstration of the developed methods utilizes hardware with several mobile phone connections from different providers. The multiplicity of connections allows data communication to be performed via a selected connection subject to the prediction. The aim is to keep the round-trip time (RTT) between sending a request and receiving the response as short as possible. A web interface (see Fig. 13) was developed for the evaluation of the AI methods. The interface visualizes the recorded data, as well as different prediction methods. Thus, the demonstration validates the success of the developed methods.

VI. FUSION AND DECISION MAKING

Employing PRYSTINE's sensors, embedded intelligence, and E/E architecture, PRYSTNE heavily focuses on the robust perception of the environment around the vehicle through the fusion of sensed data from a multitude of sensors (such as Radar, LiDAR, cameras), thus realizing PRYSTINE's FUSION technology. The integration of data and knowledge from several sources is known as "data fusion". It is a process that deals with the association, correlation and combination of data from multiple sources to achieve refined positioning and identity estimates. It is mostly used for sensor fusion since no sensor type works well for all tasks and in all conditions. Sensor-based systems generally suffer from several problems such as; "sensor deprivation", "limited spatial coverage", "limited temporal coverage", "imprecision" and "uncertainty". "Imprecision" occurs when measurements from individual sensors are limited to the precision of the employed sensing element, while "uncertainty" depends on the object being observed rather than the observing device, and occurs when features are missing, when the sensor cannot measure all relevant attributes of the perception or when the observation is ambiguous [1].

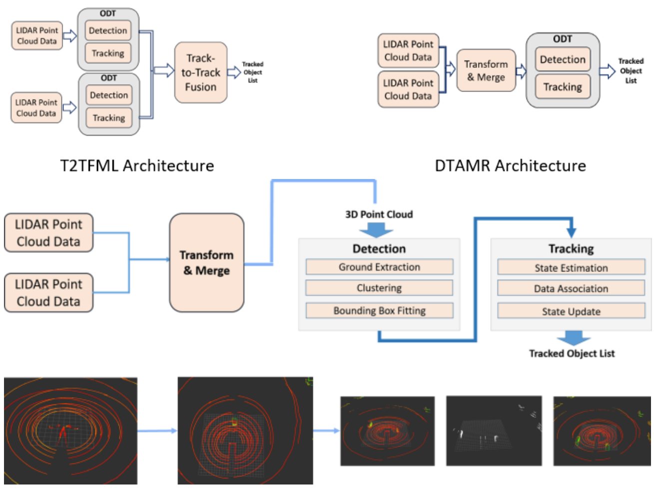

Fig. 15: Multiple 3D LiDAR fusion.

To overcome these issues, it is possible to combine data from "redundant" and "diverse" sensors to produce rich, context-aware data that eliminates the limitations in range or accuracy of the individual sensors. If data from more than one sensor or ideally more than one type of sensor is properly combined, it can be more accurate, more reliable or simply provide a better understanding of the context in which the data is gathered. Therefore, sensor fusion is extremely important for autonomous functions to improve accuracy and reduce uncertainty. PRYSTINE's vision is to implement sensor fusion, increase accuracy and derive more specific inferences from the use of a single sensor or data source alone. In line with this vision, PRYSTINE project concerns the development of the sensor fusion applications addressed to "perception", "recognition", "tracking", "estimation and control" in the following systems.

Back Maneuver Assist System: this system aims to provide an autonomous solution for backing up a truck with a trailer, which is quite difficult and prone to accidents. Most trailers are connected with a ball hitch that allows a trailer to turn along with the truck. However, this joint also makes backing up a trailer a difficult task, since the trailer can easily jut off at an angle. Additionally, it is difficult for truck drivers to recognize obstacles due to the occlusive sizes of the vehicle. Backing up a heavy-duty trailer does not come naturally and requires much practice for truck drivers, which makes the need for a back maneuver and parking assist solution essential for articulated vehicles. Towards this end, a robust solution is developed by exploiting various sensors strengthened via the power of sensor fusion.

In order to realize sensor fusion, the perception algorithms of camera, radar and LiDAR sensors are separately developed first. Objects can be simultaneously detected from all integrated sensors. As a second step, several types of sensor fusion algorithm were developed in PRYSTINE. One of PRYSTINE's sensor fusion algorithms is based on a low-level fusion of multiple LiDARs. Yet, the high-level fusion application of multiple 3D LiDAR sensors is also studied for detecting and tracking objects with high accuracy by using Track-to-Track Fusion of Multiple LiDAR (T2TFML), as shown in Fig. 14. In the high-level fusion, tracked objects from each LiDAR sensor are combined with the Covariance Intersection fusion method. Also, well-known low-level fusion algorithms of multiple 3D LiDARs were applied and tested. Furthermore, these

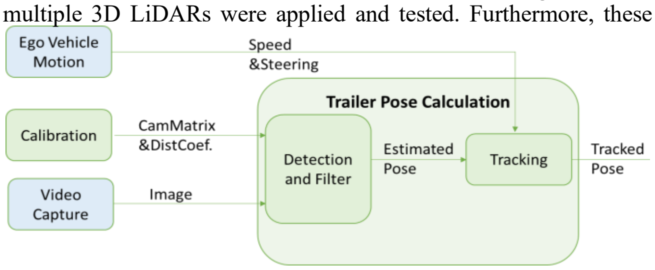

Fig. 14: Trailer Angle Detection Architecture

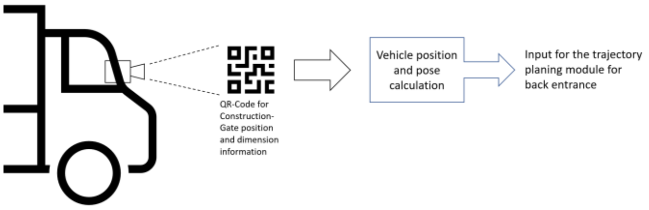

Fig. 16: Construction Gate-Recognition process-flow

two multiple LiDAR fusion strategies were compared for the first time by applying and testing them in a real vehicle [11]. This sensor fusion is based on the Detection and Tracking After Merging Raw Data (DTAMR) method as shown in Fig. 14.

Another fusion algorithm uses high-level object lists generated by the perception algorithm of camera, radar and LiDAR sensor. The intended high number of sensors was counted inside the decentralized architecture to ensure modularity, robustness and low memory requirements. The first stage in the fusion module is temporal and spatial alignment, where all objects are put to the same coordinate system. Then, the objects are associated, subsequently the state and covariance. Fusion module outputs are used to generate a global object list.

PRYSTINE's partners also developed trailer angle detection algorithm for a back parking in docking station. The relative orientation between truck and trailer is used in alignment and control mechanism. Remote sensing with a rear-view camera on the tractor and passive markers on the trailer is set up because of practicality and cost efficiency. ChArUco markers are used to detect the trailer angle. This marker is kind of binary square fiducial markers. Most significant advantage of these markers is that just a single marker can supply enough information for pose estimation. Moreover, decomposition into the marker to estimate the articulation angle, the increased number of unstable degrees of freedom, the non-linearity and unpredictability of the system make the task challenging. Therefore, the drawback of this assumption can be overcome by the fusing state observer approach with the image estimation process. The system architecture of the trailer angle detection task is presented in Fig. 15.

Furthermore, a construction gate detection algorithm was developed for back parking in construction sites. Based on the use-case requirements, it was decided to use a QR-code style sign for designating the construction gate, which is to be used for localization of the sign as well as the relative position and orientation of the truck from the sign accordingly. A barcode is a machine-readable optical label that contains information about the item to which it is attached. In practice, QR-codes often contain data for a locator, identifier, or tracker that points to a website or application. The aim here was therefore to create a construction-gate recognition solution based on a simple camera image of a QR-Code, which shall contain the specific highly accurate GNSS position of the gate position and dimensions, which is used to localize the vehicle with respect to the gate as an input for the path-planning algorithms. The process flow for the gate recognition application is shown in Fig. 16. and an example localization output

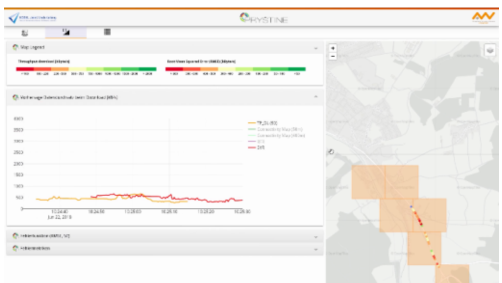

Fig. 13: Demonstrator for evaluating and visualizing the vehicle to server connection predication.

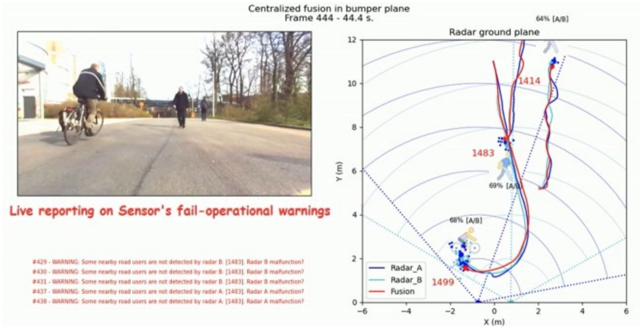

Fig. 18: VRU detection results based on Radars.

is shown in Fig. 17.

As a final study, in case of any sensor fault or data loss, sensor fusion and perception algorithms are brought to the fail operational level to achieve PRYSTINE's objective. Furthermore, a path planning algorithm was developed to extract the optimized paths for back maneuvering by the help of the dynamic map generated with using the tracked objects. These efforts were incrementally transferred to the "Heavy Duty Vehicle demonstration".

Vulnerable Road User Detection: in the EU, 22% of road fatalities are pedestrians, while 8% are cyclists. Two core technologies of PRYSTINE's FUSION technologies are dedicated to protecting these groups.

First one is focused on the Vulnerable Road User (VRU) detection based on Radars, see also Fig. 18. A key ingredient for AD is the perception of the environment around the vehicle. Such perception can guide vehicles safely and efficiently through any traffic scenario and even predict potential collision or traffic safety risks. Perception systems typically comprise intelligent sensors, compute platforms and software algorithms to detect, track and classify the objects around the vehicle. The perception quality however relies on the performance of the entire perception chain of each sensor (hardware and software). In practice, dirt or lighting may easily obstruct the sensing performance of some sensors, compute resource limitations can hamper proper data fusion, and latency may cause deadlock and inaccurate interpretations. This may lead to wrong interpretations of the environment and possibly disastrous actions of the vehicle. To avoid this, PRYSTINE developed a technology that labels the data integrity along the entire perception chain and minimizes the impact of reduced data integrity. This technology has been implemented in a multi-radar perception system and validated its operation for monitoring multiple Vulnerable Road Users (VRU).

With the technology deactivated, severe degradations of the VRU detection correctness are observed if sensors operate sub-optimally. However, activation of the technology added a data integrity indication across each data flow. Exploitation of this data integrity indictor to perform a weighted data fusion prioritized trustful data and massively boosted the perception correctness. As such, the developed technology severely contributes to the environment perception quality in realistic cases where sensors or computation resources may operate sub-optimally due to practical circumstances (e.g. dirt on sensor) or electronic health conditions.

A second VRU detection and classification solution is based on multiple vision sensors and deep learning methods. A solution for detection and classification of VRUs using You Only Look Once (YOLO) and CenterNet algorithms is developed after observing different machine learning algorithms in terms of robustness and precision. Training and inclusion of the new YOLOv4 models are completed. The YOLOv4 model is enhanced with the new activation functions like the self-regularized non-monotonic function (MISH) and its self-gated counterpart (SWISH), see also Fig. 19.

Additionally, some trials were carried out including recent data augmentation techniques (mosaic and cutmix) and some grid size configurations, with cumulative improvements over the previous

Fig. 19: Processed images from BDD100K test dataset with YOLO: original MS-COCO model from the author (left column), the BDD100K 10-class training (central column) and the BDD100K VRU-class training (right column).

results, comprising different performance-throughput trade-offs. Improvement of processes lightening the loading and processing of images on the DrivePX Board are completed.

See-Through Sight (CiThruS): the CiThruS framework is designed to virtually eliminate blind spots in the driver's field of view and thereby improve driver's awareness of occluded road users. In practice, the driver's vision is augmented with state-of-the-art vision-based technologies and the enhanced local view(s) are displayed on the windscreen or other in-cabin screen.

Two different transparency scenarios have primarily been addressed:

1) Making the surroundings of the driver's own vehicle virtually visible by mounting multiple cameras around the vehicle body, stitching the perspective corrected camera feeds together, and providing a driver with an aerial 360-degree view around the vehicle in real time.

2) Making occluded regions behind other vehicles virtually visible on the windscreen by means of vehicle-to-vehicle (V2V) video feeds from which see-through vehicle objects are generated with view synthesis and virtually immersed on top of the windscreen view.

These two scenarios have already been virtually validated with proof-of-concept demonstrations in the open-source CiThruS simulation environment [12]. Next, the corresponding experiments will be carried out in a laboratory environment by mounting cameras and CiThruS on-board processing platform on top of a mobile robot controlled with a VR headset and a mobile phone. V2V feeds are simulated with video streams sent over a wireless link from other moving devices in the laboratory.

The CiThruS framework also contains an automatic smartphone-based safety alert system that reduces road accidents and improves driver proactivity and situation awareness. The implemented safety

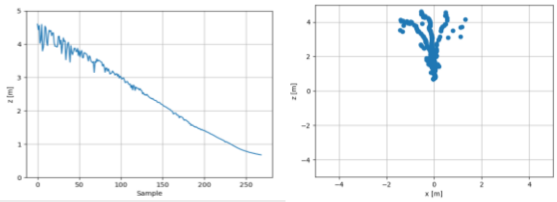

Fig. 17: Planar position & distance from the QR-Code plane

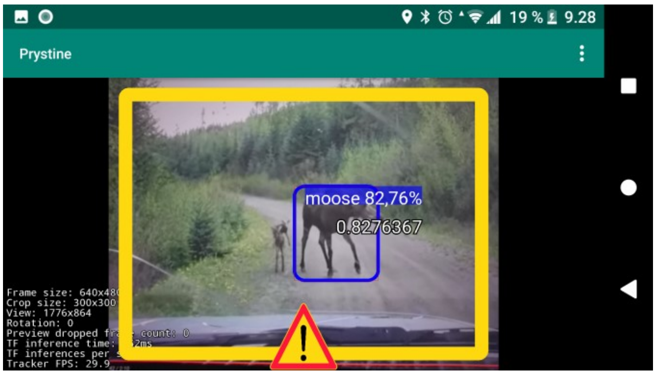

Fig. 20: CiThruS road safety alert user interface.

alert system uses AI-powered object recognition for detecting possible hazardous objects that could cause dangerous situations on road. The alerts are distributed between vehicles by using fast and low latency technologies such as MQTT.

The road safety alert system (see Fig. 20) integrates with CiThruS edge cloud computing platform by distributing all message broker activities to edge cloud computing facilities.

Vehicle Control and Trajectory Planning: a trajectory planning and control algorithm has been developed, based on a Model Predictive Control (MPC) approach to work in different road scenarios.

The algorithm allows the accomplishment of (i) way-point tracking, (ii) lane center tracking, (iii) obstacle avoidance (for fixed and moving obstacles) and (iv) constraint satisfaction (e.g., road boundaries, speed limits). In the project's first year, the MPC trajectory planning and control algorithm was designed and tested in simulation using the Matlab/Simulink simulator. In the second year, vehicle dynamics were added to the simulations and more realistic and complex road scenarios were considered in the tests, including (i) collision avoidance with one or more, fixed or moving obstacles, (ii) lane changes, (iii) emergency stops, (iv) overtaking in different conditions. In the third year, the algorithm has been integrated with other software components and tested on an embedded processor through hardware-in-the-loop (HIL) tests.

In order to proceed with hardware implementation, the algorithm is converted from Matlab/Simulink code to C++ code by means of the automatic code generation tool. Then, the C++ code is integrated with the MASA. Finally, the C++ code was deployed to a Nvidia Jetson TX2 development board. This board is equivalent to the PRYSTINE target board Drive PX2, to be used in production vehicles. HIL tests are carried out to validate the algorithm, where the vehicle is simulated on the PC, using either the Dynacar simulator provided by Tecnalia or the Matlab vehicle dynamics blockset, while the trajectory planning and control algorithm was running on the board. The HIL tests showed that the implementation of the algorithm on the embedded board and the original version of the algorithm in Matlab code provide the same

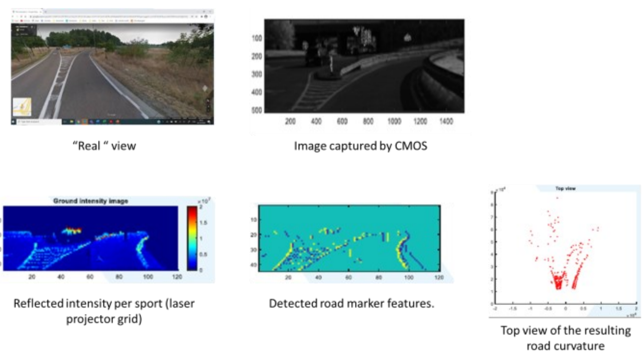

Fig. 22: Methodology to detect road curvature on demo vehicle.

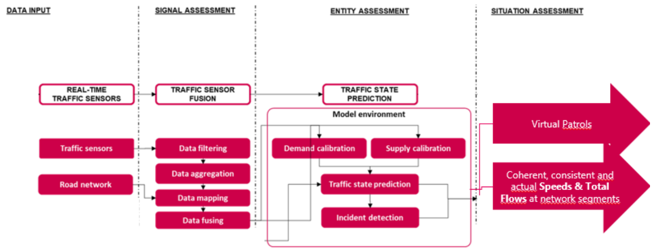

Fig. 23: Modelling architecture RTP system.

results. They also showed that the time required by the board for the computation of a single optimal trajectory and the related command inputs is less than 30 ms, allowing the application of the control command using an adequate sampling time (typically, a sampling time of 50 ms allows an effective vehicle dynamics control).

Suspension Control: a pre-emptive suspension control will improve passenger comfort. In this system, lateral and vertical road features, shapes (road curvature) and surface (roughness) can be determined. This is done through different technologies, i.e. optical image, LiDAR image and RADAR image overlapping in the ROI. The online intensity and geometry mapping generate inputs over the full working range of the suspension. This is performed through multiple gradient based methods running in parallel fused together and combined with vehicle data. One example is showed in Fig. 21 in order to detect a road curvature in the demo car.

LiDAR is used to detect each spot, and then each road marker feature is analyzed thought gradient orientation (yellow and blue). To mark an edge, two gradients, one in each direction, are needed. This method allows the identification and characterization of road specific elements of interest for the suspension system to pro-act upon. In addition, fail-operational system behavior through smart redundancy is performed, like sensing technology, parameters of interest and processing methodology.

Traffic Prediction: environment detection is a critical and essential factor for safe driving operations. The vehicle itself will be responsible for sensing its surroundings. Therefore, Radar, LiDAR and cameras need to be integrated in the vehicle. However, beyond the sensing range of these sensors, the vehicle has no real time perception of the traffic conditions. PRYSTINE's Real time Traffic state Prediction (RTP), see Fig. 22, is used as an external sensor to extend the sensing range for tactical and strategical driving operations. The sensor generates complete, consistent and actual speeds, traffic volumes and congestion levels for the road network to be used in the Decision Making Unit. RTP consists of a set of data modules and algorithms that fuse, estimate and predict network traffic states in real-time. Real time input data consists of FCD (speeds) and roadside sensors such as inductive loops, traffic controllers and travel times derived from ANPR-cameras.

Newly developed elements in RTP are bridge openings and variable messages signs (VMS) that influence temporally the road capacity and maximum speeds. Live data feeds for bridge openings and VMS are incorporated, interpreted and translated to new parameter values in the module for traffic supply.

The propagation module is extended for using alternative routes in case of dynamic road closure or congestion. The RTP modules and

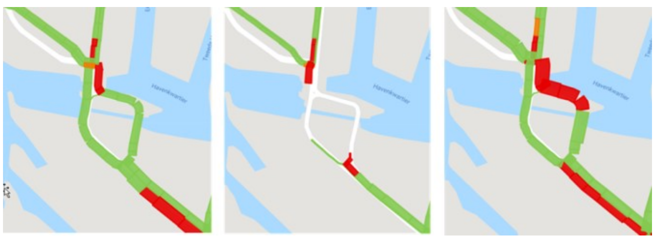

Fig. 21: Traffic state prediction before, during, and after bridge opening.

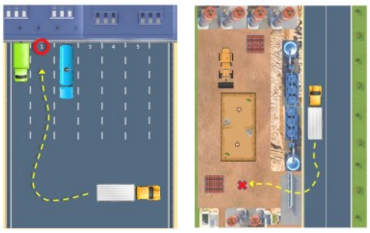

Fig. 25: FORD Demo use cases. Use Case 1: back parking in a docking station (left) and Use Case 2: back entrance to a construction site (right).

system are tested and validated for the city of Deventer (The Netherlands) and the highway A1 along Deventer (see Fig. 23). Improvements and new features are also transferred to an RTP implementation for Eindhoven-Helmond to demonstrate Shared control and arbitration applications. Testing and validation activities with the RTP system in a real-world environment have led to new ideas of application of the predictions generated in Transition of Control procedures, which will be explored in future studies.

VII. APPLICATION - HEAVY DUTY VEHICLE

One of PRYSTINE's FUSION applications focuses on realizing a heavy-duty vehicle demonstrator employing PRYSTINE's fail operational autonomous driving functions and its data fusion from a wide range of sensors (Radar, LiDAR, camera, etc.). In the context of heavy-duty vehicles, PRYSTINE advanced state-of-the-art by realizing an ambitious autonomous heavy-vehicle demonstrator for urban scenarios. Most Autonomous developments today concentrate on long-haul highway-related topics for heavy-vehicles. The need for high-precision maneuvers, combined with the size of heavy-vehicles, raises a high fatality risk, threatening the accident-free mobility scenarios in urban environments. Towards annihilating these risks, following demonstrations are pursued with corresponding use cases:

- Demo 1: a FORD heavy-duty truck (semi-size trailer),
- Demo 2: a TTS heavy-duty truck (full-size trailer).

FORD heavy-duty truck: two distinct use cases for trucks with trailers demonstrate the effective utilization of PRYSTINE's "Back Maneuver Assist System"; (i) parking in a Docking Station and (ii) backing in a Construction Site (see Fig. 24).

For both use cases, it is common that the driver needs several trials to bring the trailer in the correct position, either to dock correctly to the dedicated slot in the docking station, or to bring the truck in position on the construction site. While the main concern during the

Fig. 27: Integrated hardware components on FORD HDV.

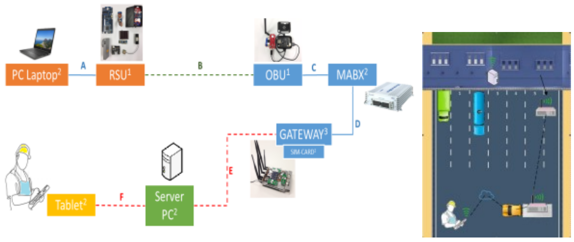

Fig. 28: Facility management hardware architecture.

docking station use case is the time spent to position the trailer, additional concerns at construction sites are the safety of the surrounding traffic and other road users, such as pedestrians. In both of the use cases, the following sensors are integrated to the demo truck; 4×LiDARs, 6×Radars, 1×Stereo-Camera and 1×Smart-Camera. Secondly, perception, sensor fusion and path planning algorithms are integrated on the target platforms tower PC and DSPACE MicroAutoBox as shown in Fig. 25.

The system provides a solution for facility managers who are in charge of a Parking Station of a warehouse or a production plant, and who need to know and track the status of the parking gates together with the trucks in operation.

In this use case, data about each parking slot of the facility is collected firstly (connection A in Fig. 26). Virtual Infrastructure and

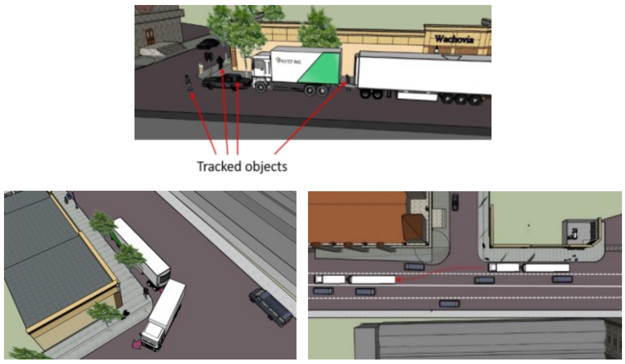

Fig. 24: TTS heavy-duty truck use-cases.

Context Sensing (VICS) system developed in PRYSTINE project merges this with environmental information sensed by the sensors on the RSU. VICS enables V2I communication and sends merged data from RSU to On Board Unit (OBU) communication module in the FORD Heavy Duty Vehicle via Bluetooth beacons (connection B in Fig. 26). The data is subsequently conveyed to the truck's central processing unit (connection C in Fig. 26). Meanwhile, vehicle-specific data is collected and slot and environment data are transmitted together with the vehicle data (connection D in Fig. 26) over the internet via LTE based connection to the Facility server using the Gateway Unit developed (connection E in Fig. 26). The Facility Server receives the parking station's slot and environment data together with vehicle data from all trucks of the facility, and

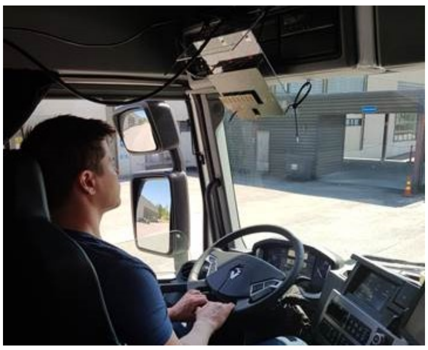

Fig. 26: 60 GHz Radar installed in front of driver's head.

Fig. 29: Micro motion data measurement with the 24 GHz radar in a passenger vehicle.

broadcasts this information over the internet (connection F in Fig. 26). Any client platform, like a tablet, displays parking slots', environment's and truck's status to the Facility Manager simultaneously.

TTS heavy-duty truck: TTS heavy-duty truck demonstration combines fail-operational ADAS functionality, decision-making system and cybersecurity solutions for safe, secure and reliable urban operations. Three use cases demonstrate the effectiveness of the developed ADAS in the urban driving scenario; i) safe assisted departure from a stationary position and safe stop when the driver is incapacitated/distracted, ii) safe right turn in urban traffic environment and iii) safe lane change (see Fig. 27).

Sensor fusion, path planning and decision making algorithms along with security solutions and sensors' data trustworthiness evaluation methods were developed and implemented in the controller platform. In the demo truck, 6 radars, 2 LiDARs, 1 smart-camera, 4 cameras for 360 degree view, and 1 stereo camera were integrated. The trust model takes advantage of supplementary weather and other relevant sensors to acquire additional information about the weather and surrounding objects, which may not be detected by the deployed main sensors (radar, stereo camera and LiDAR). In addition, a middleware for detecting HW/SW faults and security threats was designed and deployed in the controller platform. To protect the controller platform from security threats on its peripheral, an intrusion detection system was developed and integrated as part of the middleware. In addition to the surroundings perception applications, radar-based sensing can be utilized inside the cabin of a vehicle to monitor the driver's state of alertness and vital signs with a much higher level of privacy than a camera-based system. A 60 GHz SiGe based modular frequency domain imaging MIMO radar system was adapted to this purpose and tested in real driving conditions (see Fig. 28). From the radar image data, the driver's respiration rate, heart rate and heart rate variability (HRV) have been extracted. The radar is capable of monitoring several persons at the same time and track the motion of the driver. Adaptive operation using motion tracking with an optional camera system can be used to improve the reliability of measurements, especially the demanding HRV.

A more compact and low power consumption 24 GHz version was also tested, see Fig. 29. It is more robust for vital signs detection than the higher frequency version and fits well into a passenger car, but is limited to single person monitoring and has no imaging capability. The radar technology has proven to be very effective for driver monitoring application where the driver is relatively still.

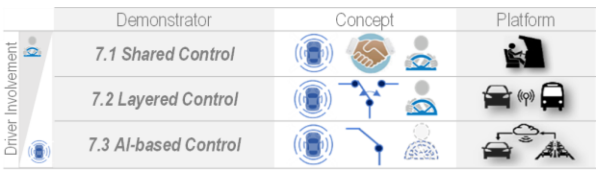

Fig. 33: Shared control and arbitration demonstrators.

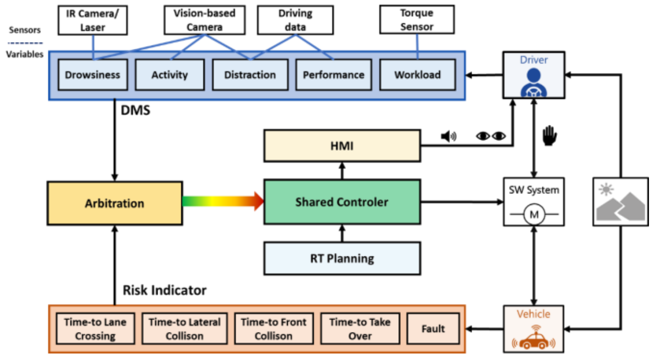

Fig. 30: Shared control and arbitration module.

VIII. Shared Control and Arbitration Applications using FUSION

Even if early efforts on AD envisioned highly autonomous vehicles with little to none driver participation, it is now clear that technical and legal issues will delay that reality. It is also true that we like to drive, and for many reasons want to "keep in control". Thus, today much effort is being focused on combining the strengths of the human driver (reasoning and deduction) and of the automation (precision and reaction time) towards improving safety and comfort on upcoming AD functionalities [13]. PRYSTINE focusses on a human-centered automated vehicle, where AD functionalities with different levels of driver's intervention are tested through shared control strategies.

Three distinct demonstrators showcase developments on providing an intelligent co-driver able to assist the driver in manual and automated mode, through safe decision-making and transitions, considering driver state awareness combined with environment and system awareness, for various levels of driver involvement (DI) (Fig. 30).

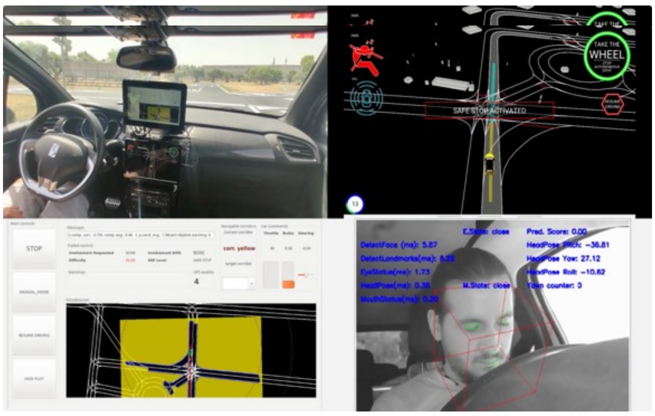

Fig. 32: Traded control and DMS validation on an automated car

Fig. 35: Shared control and arbitration simulators.

Shared Control Demonstrator: this demonstrator focuses on an intelligent co-driver with the adaptive level of assistance based on driver-automation status and different risk indicators that are present during the driving task. Until now, developments yielded the integration of different modules for achieving a driver-automation collaborative framework to share the control of the driving task in real time. Five modules are integrated: Shared Controller, Arbitration decision-making system, Driver Monitoring System (DMS), Real-Time Planner, and visual Human-Machine Interface (HMI) (

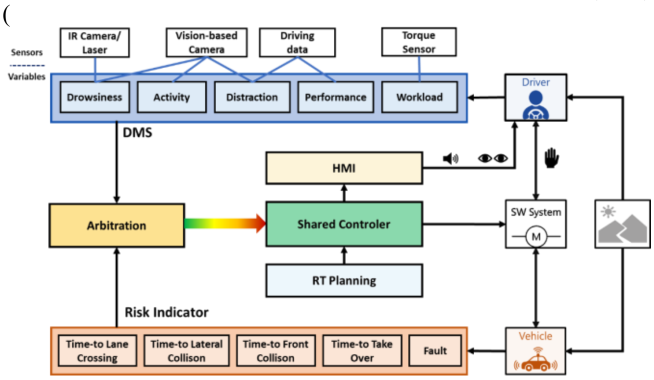

Fig. 31).

Final tests and demonstrations for this demo will take place in two facilities. Both simulators feature PRYSTINE's shared control and arbitration, real time path planning, visual HMI and DMS systems, though using different hardware and addressing different use cases. A total of four DMS approaches are implemented in two different combinations, one dealing with driver distraction use cases and the other one with driver drowsiness use cases. One DMS uses a laser and a camera to track breast motion and estimate breathing; another one uses cameras to estimate distraction and hands position; a third one evaluates driving patterns; and the fourth one uses cameras to estimate the head direction and eyes blinking. The combined DMS outputs, the risk indicators, the real-time planer and the arbitration modules define whether the torque assistance controller operates in manual, guiding or overriding mode, while the visual HMI provides information to the driver (Fig. 33).

Traded control Demonstrator: all the involved components in the demonstrator were successfully integrated and validated in simulation (ScanerStudio) and partially on a real automated vehicle in a closed driving environment (see Fig. 32). The developed architecture has been conceived to decide which is the most suitable implication level for a driver in each ODD, and act accordingly. To that end, 2 different Driver Monitoring Systems (DMS) technologies have been integrated: one relying on computer vision and AI for the eye blinking and yawning frequency estimation; another one using a combined camera-infrared sensing suite to estimate drowsy states by observing breath patterns. With this information in hand and a permanent estimation of the driving scene complexity, autonomous navigation and interaction decisions are generated, following customizable driving patterns (see Fig. 34). As

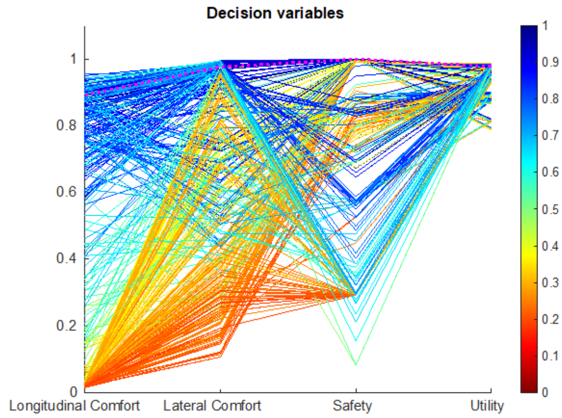

Fig. 34: Decision variables prioritization.

this complexity estimation relies on a reliable situation understanding, enhanced V2V-based perception is also being integrated to have a wider and more consistent field of view. Safety, comfort, efficiency and their related fail-operational considerations are being exhaustively tested in 6 different Use Cases, ranging from overtaking to roundabout or intersection crossings.

AI based Decision Making Demonstrator: in PRYSTINE's final year, all systems were integrated on the Demonstrator vehicle to be able to collect the required data. Several components were replaced with newer status in order to reflect the latest developments of all partners, such as LiDAR, Radar, GPS (tracker), communication equipment for traffic state prediction. When everything was up and running, the car drove multiple rounds corresponding with the chosen use cases of year 1 to collect data; which is dynamic information about other traffic (cars and cyclists). This collected data is then processed in order to be able to train the AI-algorithms which are needed for the Decision Making module. First, it was started with the algorithm for the Highway use case, which means that the algorithm learns to mimic the behavior which the driver actually performed during the data collection phase. Currently, these activities are carried out for the Urban use case as well. Next step is to validate the AI behavior which the Decision Maker shows after the training phase. This is done by replaying the use case scenarios in a simulation environment (Carla). Finally, a video is being created to showcase the project results in a movie form.

IX. Conclusion

The automation of vehicles has been identified as one major enabler to master the Grand Societal Challenges 'Individual Mobility' and 'Energy Efficiency'. Highly Automated Driving functions (ADF) are one major step to be taken. However, in order to achieve ADFs, fail-operational behavior is essential in the sense, plan, and act stages of the automation chain. PRYSTINE's target is to realize Fail-operational Urban Surround perceptION (FUSION), which is based on robust Radar and LiDAR sensor fusion, and control functions in order to enable safe AD in urban and rural environments. This work highlights the visions of PRYSTINE's research and development activities and showcases some groundbreaking results achieved until PRYSTINE's third year.

Acknowledgment

The authors would like to thank all national funding authorities and the ECSEL Joint Undertaking, which funded the PRYSTINE project under the grant agreement n° 783190.

References

[1] N. Druml et al., "Programmable Systems for Intelligence in Automobiles (PRYSTINE): Technical Progress after Year 2," 2020 23rd Euromicro Conference on Digital System Design (DSD), Kranj, Slovenia, 2020, pp. 360-369.

[2] A Dubey, A Santra, J Fuchs, M Lübke, R Weigel, and F Lurz. Integrated classification and localization of targets using Bayesian framework in automotive radars. 2021 IEEE International Conference on Acoustics, Speech and Signal Processing, 6-11 June 2021, Toronto, Canada.

[3] N. P. Jouppi et al., "In-datacenter performance analysis of a tensor processing unit," 2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA), Toronto, ON, 2017, pp. 1-12. DOI: 10.1145/3079856.3080246.

[4] Y.-H. Chen, T.-J Yang, J. Emer, V. Sze, "Eyeriss v2: A Flexible Accelerator for Emerging Deep Neural Networks on Mobile Devices," IEEE Journal on Emerging and Selected Topics in Circuits and Systems (JETCAS), Vol. 9, No. 2, pp. 292-308, June 2019.

[5] T. Moreau et al., "A Hardware–Software Blueprint for Flexible Deep Learning Specialization," in IEEE Micro, vol. 39, no. 5, pp. 8-16, 1 Sept.-Oct. 2019, DOI: 10.1109/MM.2019.2928962.

[6] Novickis, R., Levinskis, A., Kadiķis, R., Feščenko. V., Ozols, K. (2020). Functional architecture for autonomous driving and its implementation. 17th Biennial Baltic Electronics Conference (BEC2020), Tallinn, Estonia.

[7] Justs, D., Novickis, R., Ozols, K., Greitāns M. (2020). Bird's-eye view image acquisition from simulated scenes using geometric inverse perspective mapping. 17th Biennial Baltic Electronics Conference (BEC2020), Tallinn, Estonia.

[8] N. Druml et al., "PRYSTINE - Technical Progress After Year 1," 2019 22nd Euromicro Conference on Digital System Design (DSD), Kallithea, Greece, 2019, pp. 389-398.

[9] Macher G., Druml N., Veledar O., Reckenzaun J. (2019) Safety and Security Aspects of Fail-Operational Urban Surround perceptION (FUSION). In: Papadopoulos Y., Aslansefat K., Katsaros P., Bozzano M. (eds) Model-Based Safety and Assessment. IMBSA 2019. Lecture Notes in Computer Science, vol 11842. Springer, Cham.

[10] Naveen Mohan, Martin Törngren, " AD-EYE: A Co-Simulation Platform for Early Verification of Functional Safety Concepts.", SAE Technical Paper 19AE-0203/2019-01-0126

[11] E. A. Tasdelen, V. Sezer, Comparison and Application of Multiple 3D LIDAR Fusion Methods for Object Detection and Tracking, 2020.

[12] T. Niemirepo, J. Toivonen, M. Viitanen, and J. Vanne, "Open-source CiThruS simulation environment for real-time 360-degree traffic imaging," in Proc. IEEE Int. Conf. Connected Vehicles and Expo, Graz, Austria, Nov. 2019.

[13] Marcano, Mauricio, et al. "A Review of Shared Control for Automated Vehicles: Theory and Applications." IEEE Transactions on Human-Machine Systems, 2020