UP-Drive

Automated Urban Parking and Driving

Research and Innovation Action - Horizon 2020
Grant Agreement Number 688652
Start date of the project: 01-01-2016, Duration: 48 months

Deliverable D 2.2

First vehicle platform fully operational

Status: Revision 1.0

Lead contractor for this deliverable: VW
Due date of deliverable: 31.12.2016
Coordinator: VW

Project co-funded by the European Commission within HORIZON 2020		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
SEAM	Restricted to partners of the SEAM Cluster (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Executive Summary

This deliverable documents the functionality of the first vehicle platform. It provides a thorough analysis of the sensor setup, presents the high-level processing framework, reports on communication capabilities and provides a brief overview of the safety elements and policies.

Contributing Partners

	Company/Organisation	Name
Document Manager	VW	Wojciech Derendarz
Partner	CVUT	Radim Šára
Partner	CVUT	Martin Matoušek

Revision Table

Document version	Date	Notes
1.0	29.06.2017	document ready for submission

Contents

1 Sensor data integrity validation

UP-Drive vehicle collects the following sensor data:

- 360° laser scanners: point clouds
- 4 layer laser scanners: point clouds
- Long range radar: objects and clusters
- Short range radars: clusters
- TopView cameras: images
- Trifocal camera: images and objects and other semantic data like traffic lights
- Odometry: relative vehicle motion
- Consumer grade GPS: coarse position and orientation data
- Consumer grade IMU: relative vehicle motion
- Ground truth DGPS: full 6DOF position and motion data

The data availability and format have been already reported in Deliverable 2.1. Here we analyze the integrity of the available sensor data. This is done in 3 steps. First, the measurements (i.e. the content of the data streams) are analyzed. Second, the calibration data is verified. Finally, the time-behaviour and timestamp information is analyzed.

1.1 Sensor measurements verification

1.1.1 Distance / 3D measurements

In this section the integrity of the 3D measurements of our sensor system is analyzed. The analysis has been performed by observation in real life scenarios and performing additional controlled experiments. In general it can be said, that the sensors perform in accordance to specification. Nonetheless, some limitations - arising from the fact that we are using the sensors in real (and not ideal) environments - have been observed. Those observations are reported below.

360° laser scanners

The sensors work in accordance to the specification. Range measurements are observed up to the distance of 80-100m. The main practical limitation of the sensor arises from what is also its big advantage: it provides a large vertical field of view of 30° and is thus able to perceive real 3D structures. However, as the sensor has 16 layers, the spacing between the layers amounts to about 2°. This - combined with the fact that the sensors are mounted on top of the vehicle - leads to a setup in which objects the size of a car can be "hidden" between the layers of the sensor. This effect is amplified by the fact that usually only the lower parts of a car provide stable measurement data whereas the rays hitting the windshield get reflected away. We remedy this drawback by using multiple 360° sensors and titling them relative to one another in order to obtain best possible coverage across the vertical field of view. The experiments show that with such a setup, passenger cars provide stable measurement data

up to a distance of 40-45m. This contrasts with range available for large vertical structures, which amounts to double that value and beyond.

Given those observations, we can conclude that the 360° laser scanners in our setup are very useful sensors for perceiving the vertical profile of the road or the objects - but only up to moderate distances. We believe that the vertical profile will turn out especially crucial at tackling the object classification task and the even more demanding tasks of cue detection for intention estimation of pedestrians and cyclists (body orientation, leg configuration, head pose, gestures).

4-layer laser scanners

The sensors work in accordance to the specification. Range measurements are provided up to a distance of 150-200m. Unlike the 360° laser scanners, the vertical field of view of the 4L-scanners is only about 3°. As the sensors are mounted horizontally at the height of vehicle bumper, such configuration allows for very stable object detection - at least in situations where the ego-car is level and the road-surface is flat. Deviations from that ideal configuration can lead to substantial reduction of perceivable range.

So far, that limitation has not manifested itself during real tests performed in urban environments of flat Northern Germany. Consequently the 4-layer laser scanners - with their high range, ability to perceive object contours and 360° coverage due to the 4 sensor configuration - seem to be the core sensing modality, around which the obstacle detection of UP-Drive will be built.

Long range radar

The sensor works in accordance to the specification. Range is beyond 200m. The data for static objects looks cluttered, which is typical for radar sensors. The object data for dynamic objects looks very clear though. It should be noted that radar sensors are usually only able to reconstruct the radial component of the object velocity. This means that for a front sensor the longitudinal motions get well measured, whereas velocity of crossing traffic is measured inadequately. However, in case of the long range radar from our setup - due to its internal tracking - the velocity vectors of crossing traffic get correctly reconstructed.

Short range radars

The sensors work in accordance to the specification and detect objects up to a distance of 70m. With each echo, its relative radial velocity is provided. Main drawback of the sensors is the limit on maximum number of echos provided - with strongest echos being provided and weaker suppressed. This often results in situations, where echos are concentrated on the closest objects with many echo-points per object. The detection or suppression of object at longer distance depends on the obstacle configuration at closer distance, which is quite unfortunate.

In consequence, the resulting data seems quite noisy (for any given object, sometimes echos are provided and sometimes suppressed) and not particularly reliable as single source of measurement. As the sensor provides not only distance but also velocity information, it seems very suitable to augment the object detection based on laser scanners.

Trifocal camera

The sensor works in accordance to the specification. In case of obstacle data its strengths are the angle accuracy and object class information. The object radial distance is subject to high uncertainty. Another limitation is that only front and rear faces of vehicles get provided - information on crossing traffic is not provided. Apart from obstacle data, the sensor also provides detections of lane markings or delimiters, traffic lights and signs. All those

functionalities have been validated to work as expected.

1.1.2 Ego motion measurements

In this section, the following signals are analyzed and compared:

- **GPS**: Consumer grade GPS, providing readouts at 6Hz.
- **RT3k**: A high quality inertial and GPS Navigation System with support of Real-Time Kinematics (RTK). It provides output at 100Hz and is only used as a reference system in the project.
- **Odometry**: Wheel odometry fused together with a next-generation automotive IMU. Output is provided at 100Hz.

In a first step, we compare consumer grade GPS measurements with RT3k both in WGS 84. It is expected that RT3k provides position estimates that are close to the true positions. As a sanity check, the estimated RT3k trajectory is plotted in Google Maps in Fig. 1.

Figure 1: Google maps overlayed with the estimated RT3k trajectory in Wolfsburg.

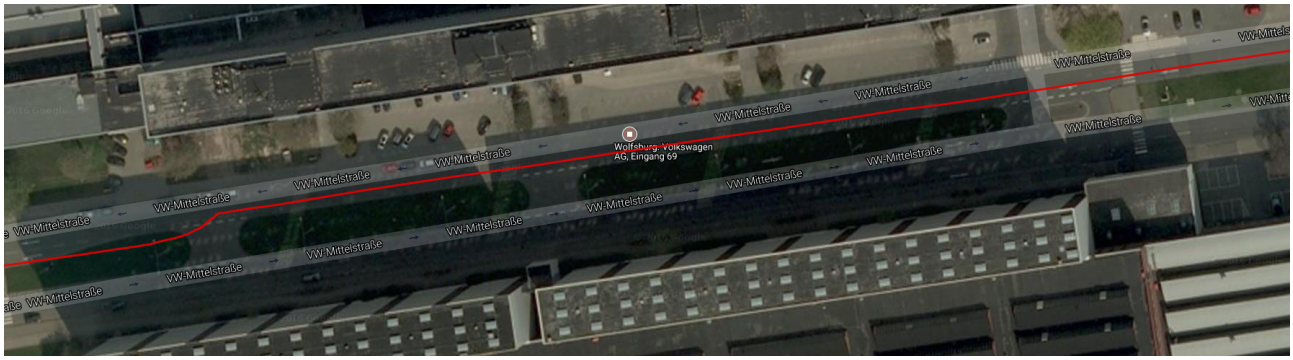

Figure 2: Magnified version of Fig. 1.

At first sight, the trajectory seems to be perfectly following the street. However, looking at the left part of the trajectory shown in Fig. 2, it seems that there can still be consistent drift for a longer section of the trajectory.

Similar observations can be made with the consumer grade GPS signal for which the same section of the trajectory as in Fig. 2 is shown in 3. Clearly, the trajectory is drifting off the street, indicating erroneous position data.

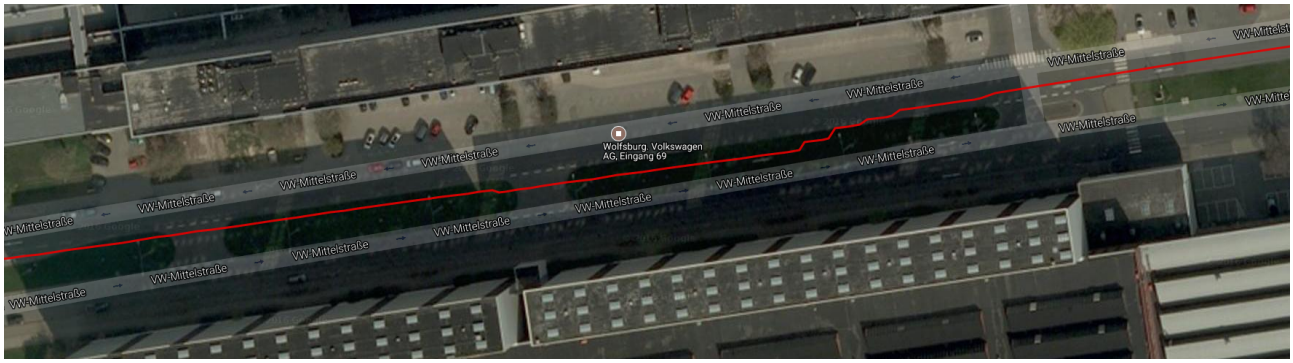

Figure 3: Consumer grade GPS trajectory of the same section as in Fig. 2.

Fig. 4 compares RT3k and consumer grade GPS and suggests that the RT3k trajectory is overall smoother. In fact, manual inspection of the data led to the conclusion, that the RT3k position data is significantly more accurate in general.

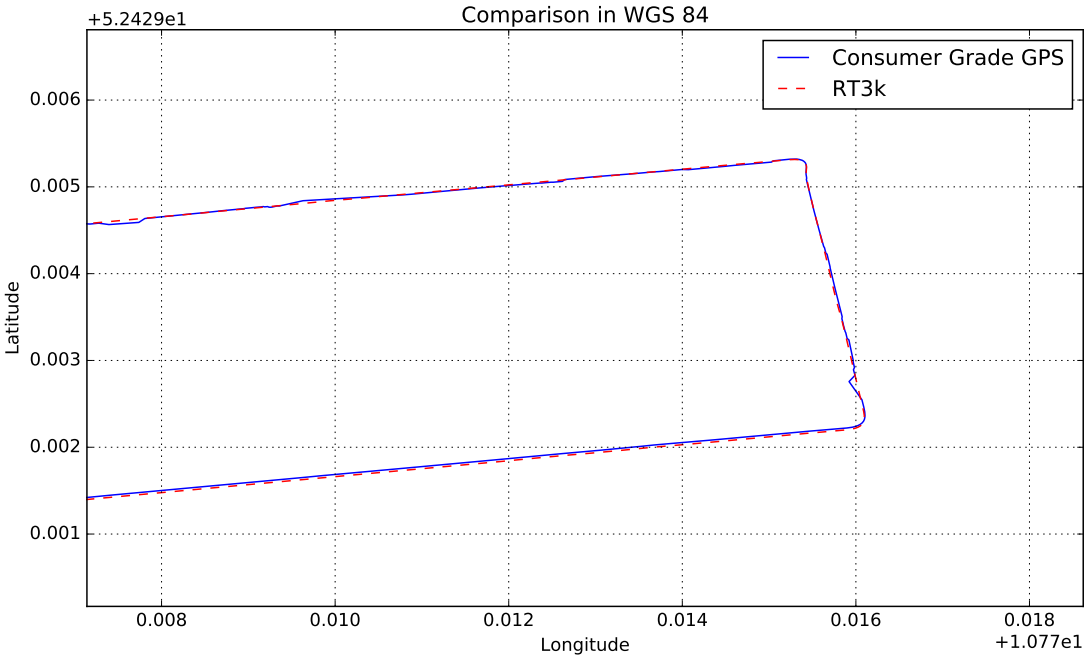

Figure 4: Consumer grade GPS trajectory compared with RT3k.

An additional source of information is the fused signal of wheel odometry and an automotive IMU - which will be stock sensor in the upcoming vehicle models. It does not make use of global position measurements and is simply called odometry for future reference. It is expected that a trajectory estimated solely with this signal will accumulate drift over time. Fig. 5 compares both the RT3k trajectory with a trajectory estimated with the wheel odometry

and the automotive IMU.
As anticipated, the odometry trajectory drifts away from the globally more accurate RT3k trajectory. It is only a coincidence that the final position of both trajectories almost coincide. Still, the scale of the odometry trajectory is close to the true scale and the angular error stays small throughout the trajectory.

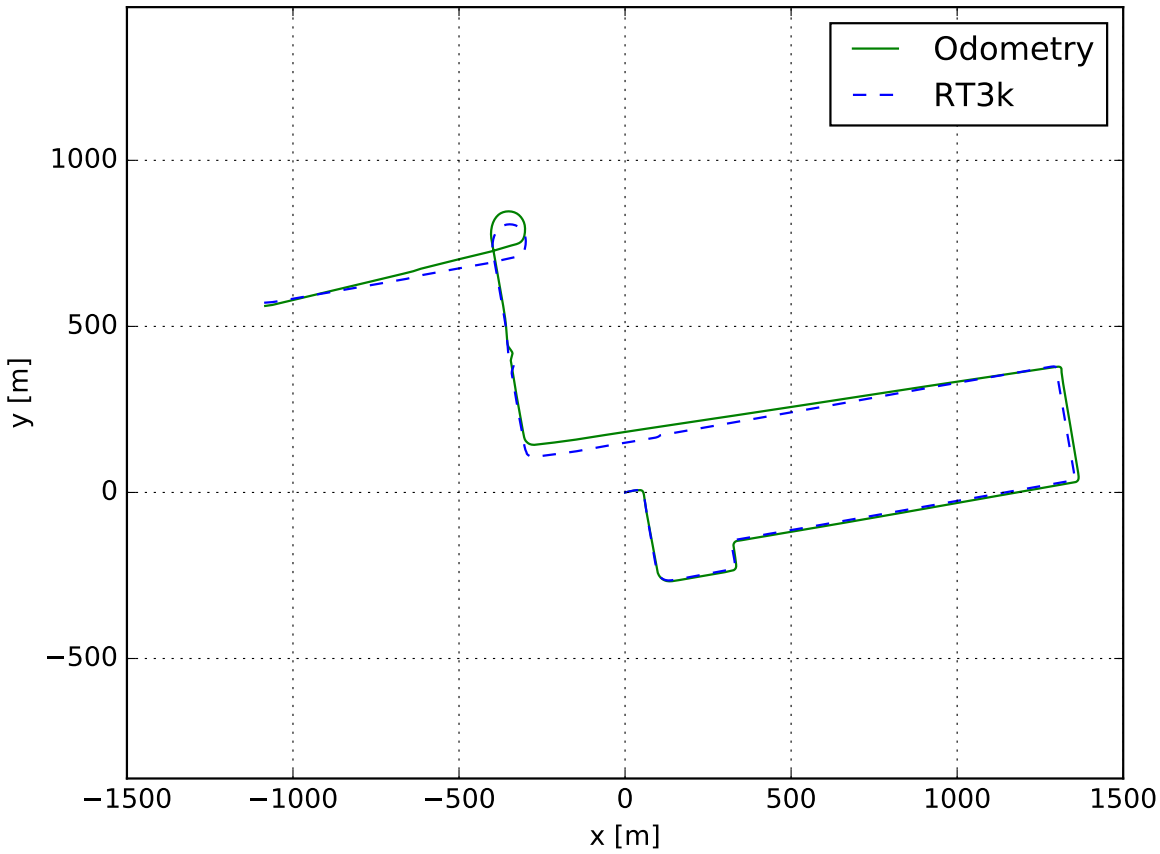

Figure 5: RT3k trajectory compared to odometry trajectory.

To summarize, the consumer grade GPS fused position estimate is surprisingly accurate overall. It proves to be a valuable source of information that could be used to aid localization. While RT3k position estimates are, as expected, more accurate than consumer grade GPS estimates, it does still have drift (usually less than one meter) at some parts of the trajectory. The odometry position estimate is more accurate than anticipated, keeping the angular error small even for trajectories of several kilometers of length.

1.2 Calibration data and its verification

1.2.1 LiDAR data

Calibration data for LiDAR scanners calibration consist of a drive into and out from a calibration room. To support 4-Layer LiDAR calibration the room is equipped with slanted planes in the corners. The room has no special properties, it represents just a convenient infrastructure that is available for repeatable calibration experiments. No assumptions are made about the room, the only requirement on its structure is that it is as simple as possible while

sufficiently rich in structure. The structure should be asymmetric to avoid multiple solutions during the automatic registration of LiDAR scans.

A single data frame consists of a single revolution of the 360° LiDAR and a half-sweep of the 4-Layer LiDAR. More details on this data is found in Deliverable D4.1. Results reported here used seven frames, four from a drive out of the calibration room and three from the drive in the calibration room. Fig. 6 shows two sub-frames from one of the four 360° LiDARs.

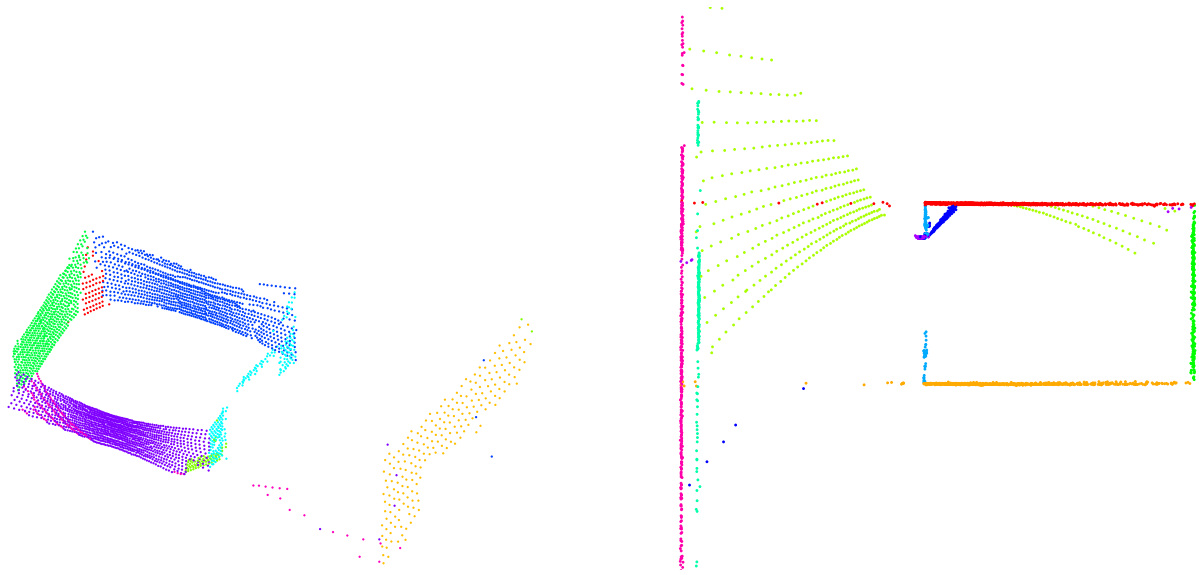

Figure 6: Segmented sub-frames from one 360° LiDAR. They come from different frames. Colors represent automatic segmentation. The car is inside the room in the left frame. An oblique view is shown of three walls of the room (blue, green, violet), the door opening (cyan), the slanted boards in the corners opposite to each other: a far segment (red) and a close segment (green, partly obscured), the distant wall visible through the room door (orange) and the floor (magenta). The car is half-way out of the room in the frame shown at right from above the vehicle. Segment colors do not correspond at this stage of processing. Point density is reduced to one-tenth of the scanner resolution. Other scans capture more of the floor, hence the floor plane (or the ceiling plane) is not underrepresented in calibration data.

Frames are processed in several stages, as detailed in Deliverable D4.1. The four 360° LiDAR sub-frames are first registered relative to each other with high accuracy, with the help of accurate planar segmentation such as the one shown in Fig. 6. This is done at one-tenth of scanner resolution, so that the sampling density is approximately uniform in all directions. There are typically 40 LiDAR-to-LiDAR plane correspondences for each LiDAR pair that are collected from all frames. These correspondences support the registration transformation between the sensor coordinate frames of individual 360° LiDARs.

In the next stage the registered plane segments from all frames are fused and the registration parameters are further refined. The fused segments are then matched to 3D positions of visual markers on the walls and the floor. This supports the registration transformation from the common 360° LiDAR coordinate frame to the room coordinate frame.

Approximate positions of 360° LiDARs in the car body coordinate frame are obtained independently by means of an external optical calibration system.[1] Poses estimated from the

[1] MoveInspect DPA system from AICON 3D System Gmbh. See http://www.aicon3d.com/fileadmin/user_upload/produkte/en/moveinspect/dpa/pdf/Brochure_MoveInsepct_DPA_EN.pdf

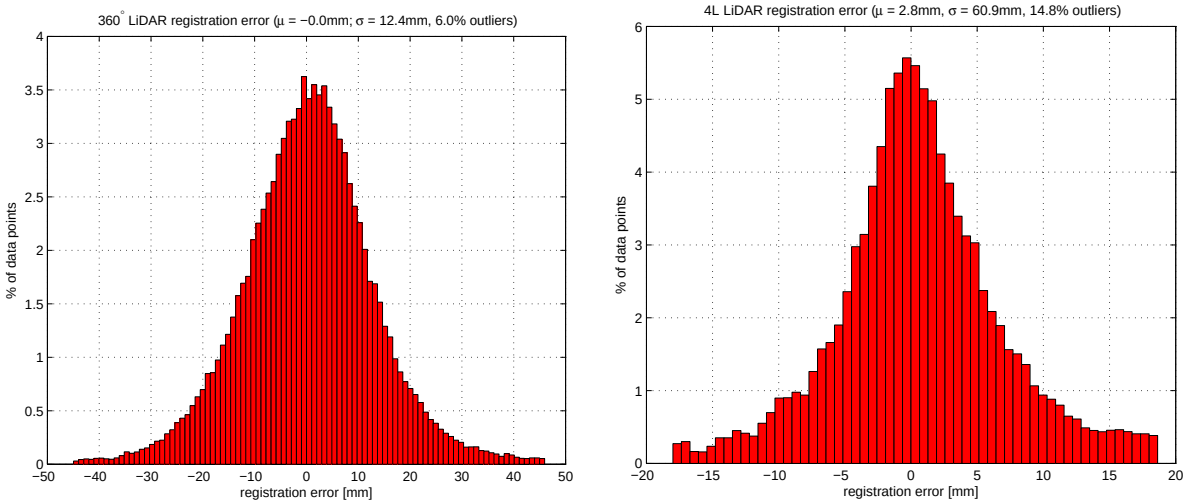

(a) 360° LiDAR (b) 4-Layer LiDAR

Figure 7: Registration residual error histograms (inliers only). (a): Residual errors for 3D point i from sensor j are the signed distances of that point to all planes from other sensors to which plane j is matched to. The error function is robust. We see a nearly normal error distribution, a 6% outlier rate, mean error under 0.0 mm and a standard deviation of 12.4 mm. These results are very good considering declared sensor accuracy of 30 mm. (b): Residual errors for the 4-Layer LiDAR indicate lower resolution which is consistent with its declared accuracy.

above 360° LiDAR registration are matched to the externally calibrated poses. The mean residual translation error is 16 ± 7 mm and the mean orientation error is $0.3° \pm 0.14°$.

Fig. 8 illustrates the result. This partly validates the automatic calibration. The side-effect of this step is a coordinate transformation from the 360° LiDAR coordinate frame to the car body coordinate frame.

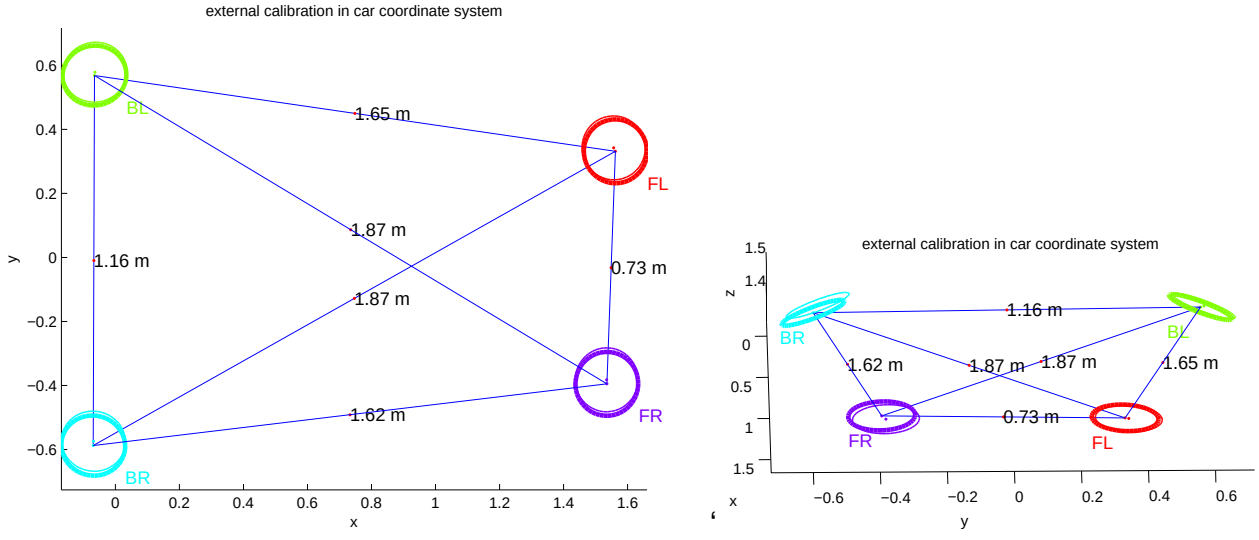

(a) top view (b) oblique front view

Figure 8: 360° LiDAR poses obtained from automatic calibration superimposed on poses obtained by means of an external calibration procedure. Numbers are measured distances. Positional deviations are indicated by small dots near the inter-linked circle centers and the angular deviations are indicated by the differences between thick and thin circles in each 360° LiDAR.

Inter-frame 360° LiDAR registration can also be used to estimate car egomotion between individual frames. Fig. 9 shows the result. This provides a further validation of the 360° LiDAR external calibration.

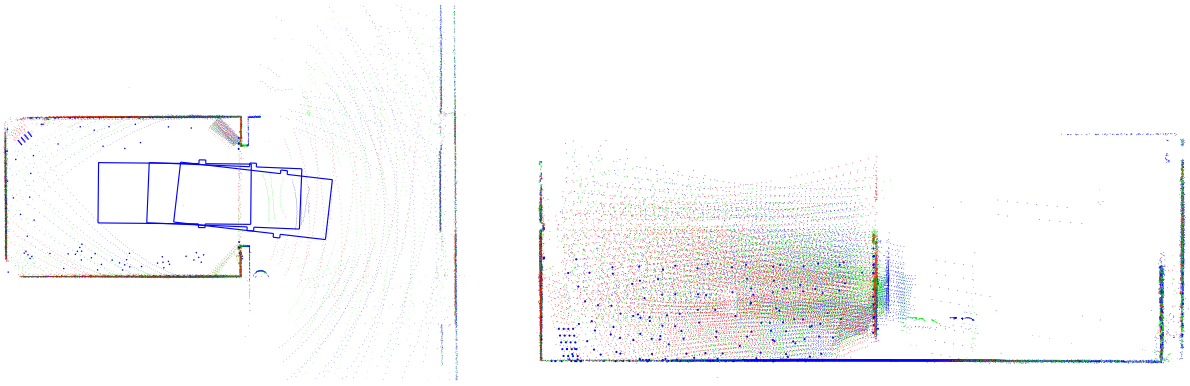

(a) top view (b) side view

Figure 9: Registered 360° LiDAR data from three frames shown in car body coordinate frame. Small dot colors represent individual scanners identity. Larger blue dots are positions of external calibration markers on the walls and the floor. Car outline is shown for the three frames (in the side view the outline coincides with the ground plane). We see that fine room features are captured without smearing, the ground plane segments in all three frames coincide, and that 3D point colors are well intermixed, which means that the registration is very accurate. This is vector graphics suitable for detailed inspection.

The 4-Layer LiDAR scans consist of three closely spaced layers, as detailed in D4.1. Relative to the size of the calibration room the scan is essentially just a single line of 3D points. This makes the problem of automatic registration difficult because it is considerably more difficult to register lines onto planes than planes onto planes. Extensive experimentation shows that the only working method for estimation of the transformation from the 360° LiDAR to the 4-Layer LiDAR coordinate frame is based on local optimization based on the same point-to-plane error model as in the 360° LiDAR. Initial estimate is provided by the external optical calibration system. The final estimate is based on a Bayesian model with a weak location and orientation prior centered at the initial parameter estimate. The result is illustrated in Fig. 10.

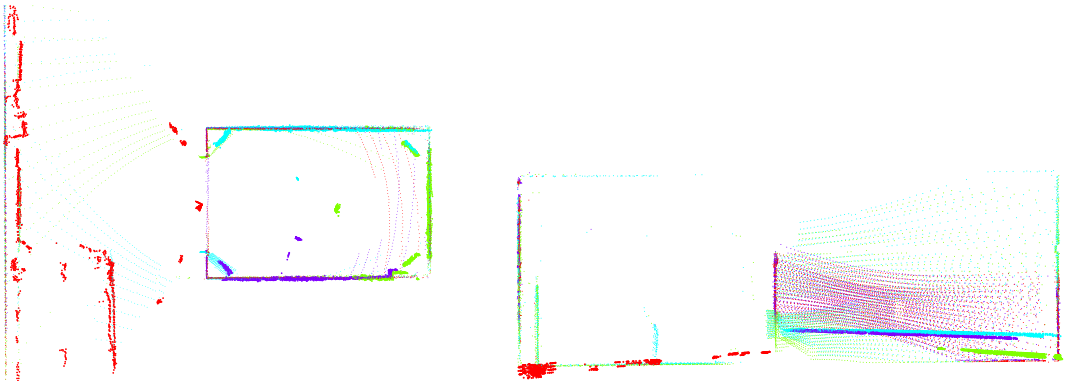

(a) top view (b) side view

Figure 10: Registered 4-Layer LiDAR data (crosses) and 360° LiDAR data (dots). This is vector graphics.

1.2.2 RADAR data

Approximate positions of RADAR sensors in the car body coordinate frame are obtained independently by means of an external optical calibration system mentioned above (we call this procedure a 'precalibration'). The accuracy of resulting poses is already within the noise level of the RADARs. To verify the calibration and provide a possible calibration refinement we tested a maximum-likelihood method.

The method calibrates the yaw angle of RADAR orientation relative to fused LiDARs. It is assumed the LiDARs are already extrinsically calibrated in the car coordinate system. Besides an extrinsic RADAR precalibration the method does not require any additional calibration aids. It uses data from infrastructure, based on a sufficiently long drive in a sufficiently complex environment (real traffic in an urban environment). Data were collected from a 3.5 minute long drive in normal inner city traffic. Fig. 11 shows four input frames from the sequence. We can see that obstacles detected by LiDARs and RADARs are already quite consistent.

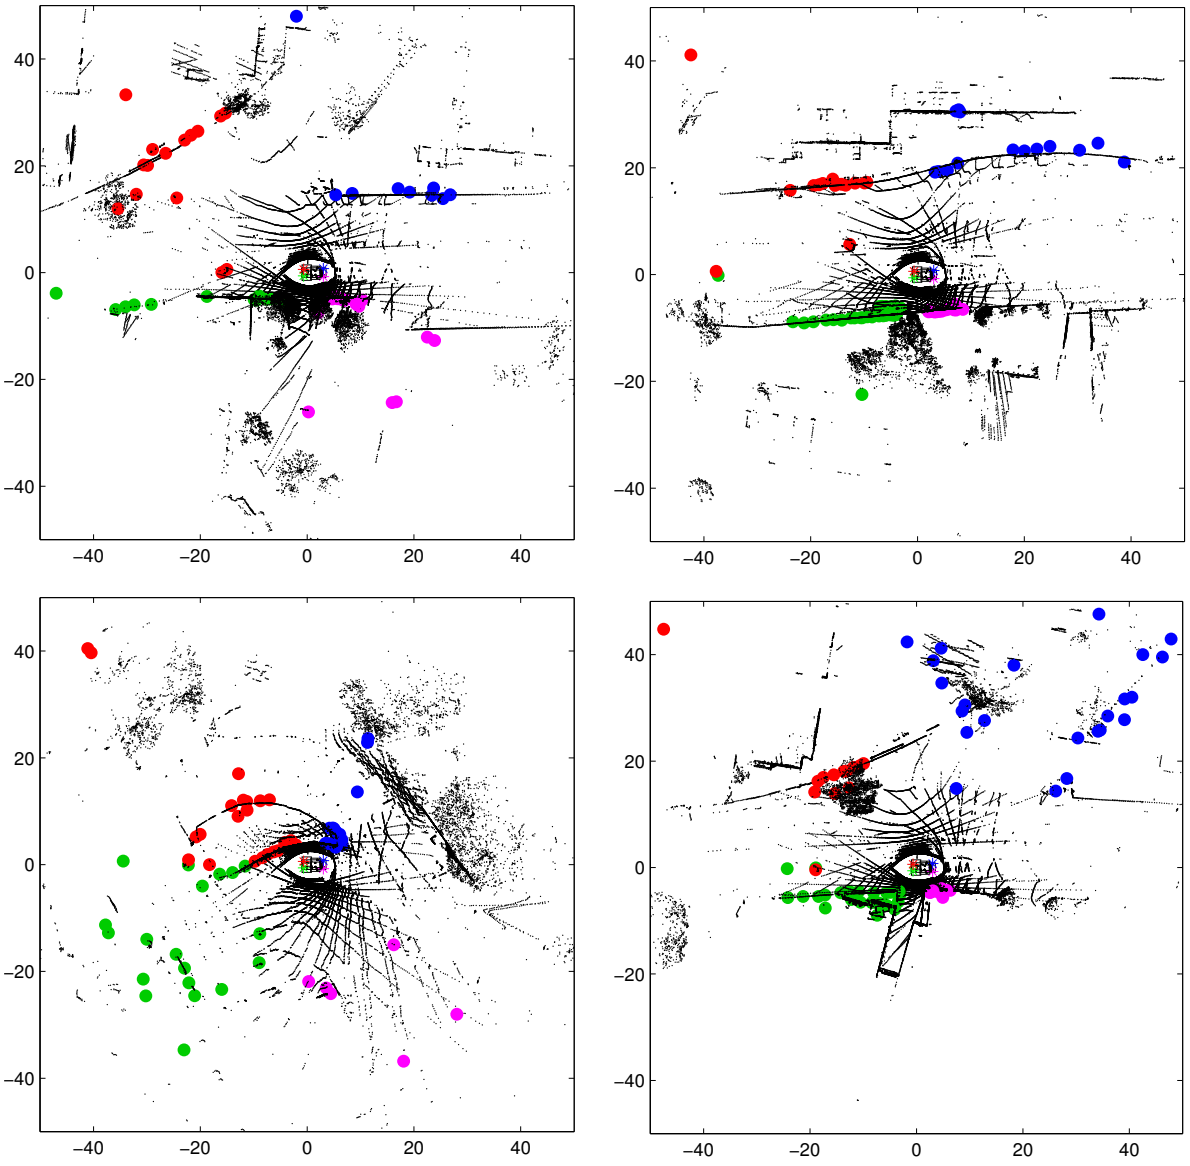

Figure 11: Pre-calibrated short range RADAR data (colors represent individual RADARs) and fused 360° LiDAR data in car coordinate frame (black). This is vector graphics suitable for detailed inspection.

Fig. 12 shows all radar data collected from the drive. This figure is interesting on its own because it shows the noise level and an uneven distribution of noisy RADAR responses. The source of this unevenness must be investigated but we believe it has little effect on the subsequent calibration refinement method.

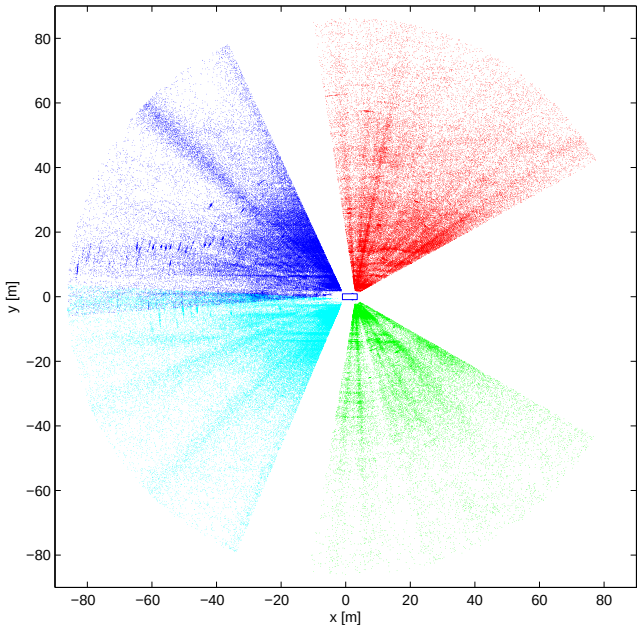

Figure 12: Distribution of all RADAR measurements in the calibration refinement drive. Here, RADARs orientation is determined from the precalibration. Scale is given in meters. This is a vector graphic.

We assume RADARs provide 2D points. We take 3D points from the LiDARs and, together with the RADAR points, we project them onto the ground plane in the car coordinate system. If the RADAR-LiDAR orientation was correct, there should be coincidence of sparse RADAR measurements and dense LIDAR measurements. The calibration verification/refinement is based on a robust semi-parametric maximum-likelihood method that maximizes the point set alignment. The method works by updating accumulated likelihood frame by frame. Its on-line nature makes it suitable for on-line calibration monitoring and/or refinement.

Fig. 13 shows calibration curves for four short-range radars mounted on the vehicle. The refined orientations of the RADARs are shown in the legend. They are also summarized in the following table:

RADAR	yaw angle refinement
front left	$-2.3 \pm 0.7°$
front right	$+2.0 \pm 0.6°$
back left	$+0.6 \pm 0.4°$
back right	$+0.2 \pm 0,2°$

The corrections do not make visible difference on data frames shown in Fig. 11. But the curves in Fig. 13 suggest that it is relatively easy to obtain accurate yaw calibration with even a very approximate precalibration because the maxima are well-defined in the angular interval of $\pm 10°$ (and possibly larger). The calibration curve peaks appear well-defined already after about 90 sec of driving and remain quite stable, with their precise positions and accuracies slowly improving over time.

Interestingly, concentration of the curves is lower in the front two RADARs in Fig. 13. Whether this is related to the particular dataset choice, individual sensor characteristics, or uncompensated motion of LiDAR data remains unclear and must be investigated.

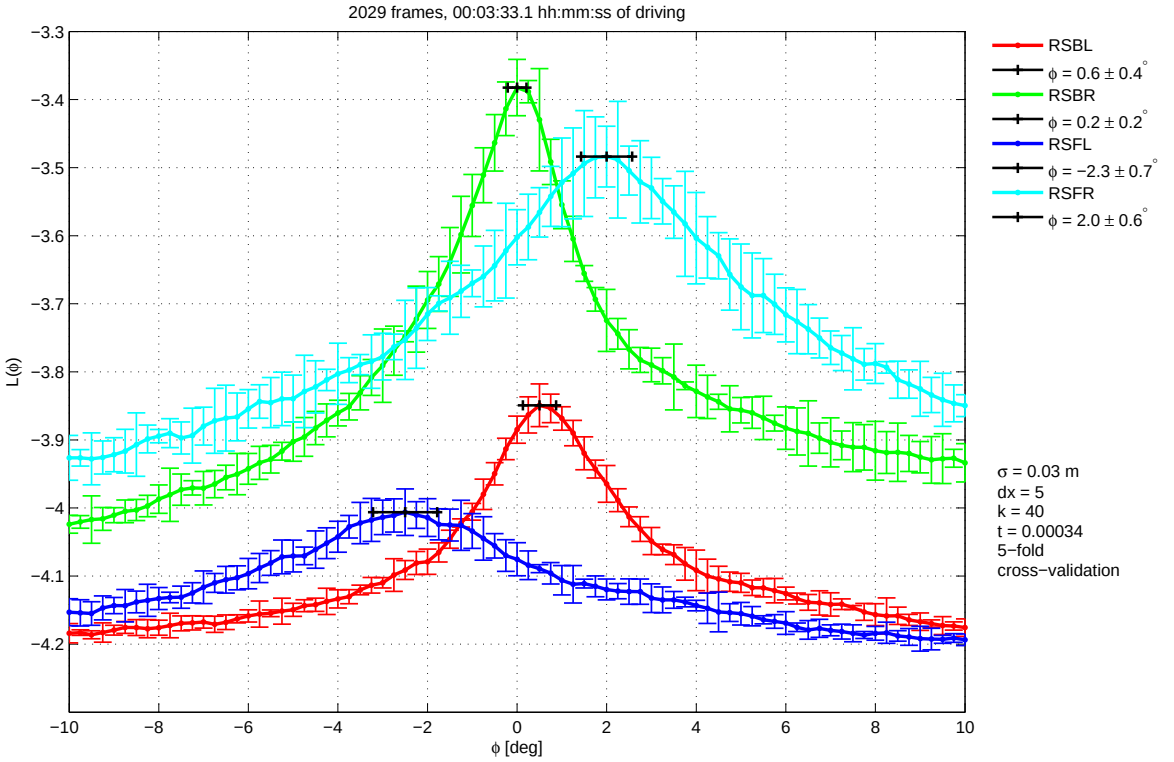

Figure 13: Calibration curves (likelihood function) for four short-range RADARs (RSBL is the back left RADAR, RSBR is the back right RADAR, RSFL, RSFR are the left and right front RADARs, respectively). The curves are mean likelihoods and their standard deviation intervals over five random subsets of data. The x-axis is the yaw correction angle in angular degrees. Peaks are within 2-degree tolerance from the precalibration. Black horizontal bars illustrate the mean and accuracy of their location. Refined yaw angle orientations are shown in the legend and are summarized in the text.

1.2.3 LiDAR to Camera

Cameras are calibrated extrinsically by resection from a set of coded calibration markers on the walls of the calibration room. The initial estimate is then refined by a bundle adjustment method. This is a standard procedure. The resulting extrinsic parameters are expressed in the coordinate frame of the calibration markers.

Since the walls of the calibration room are simultaneously measured by LiDARs, it is possible to compute the transformation from the calibration marker coordinate frame to the car coordinate frame, as long as the LiDAR points are already expressed in the car frame by means of the procedure described in Sec. 1.2.1. The transformation is estimated from several car positions in the calibration room that are used for the camera external calibration, which improves accuracy.

Fig. 14 shows projections of all fused LiDAR points to a TopView camera. We can see good accuracy of the projection: Thin objects match perfectly (cf. cables hanging from the ceiling), object edges also match in both modalities. The upper right frame shows a slight misalignment on the two persons moving in front of the car, this is due to their motion and the time latency between the 360° LiDAR and 4-Layer LiDAR. The bottom two frames have a car

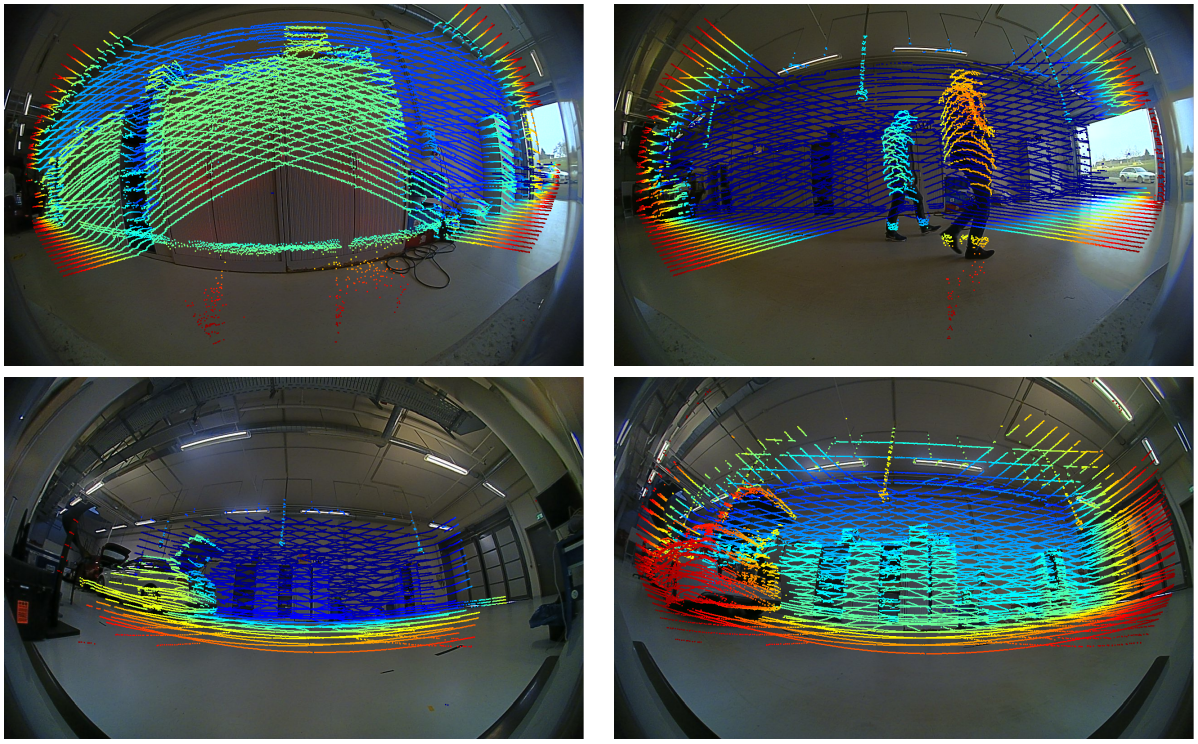

Figure 14: Several frames of fused LiDAR points (both 360° LiDAR and 4-Layer LiDAR) projected to one of the TopView cameras.

on the left side of the scene, where the LiDARs point projections might create an impression of inconsistence but in fact the LiDAR points are at their proper 3D positions, on the car and behind it. What we also see in the images is the consistency of depth estimates from both LiDARs: The colors blend seamlessly and we cannot easily distinguish which line of 3D points originated from which LiDAR.

1.3 Time behaviour and time-stamp analysis

Time stability an synchronicity of sensors is analysed using sample time-stamps of data packets from sensors. The 360° LiDAR sensors and GPS based ego-motion sensors assigns to the data the time-stamps derived from GPS time. The time-stamps for other sensors are derived from system clock at the time of arrival to the acquiring PC.

1.3.1 Period of time-stamps

Period between time-stamps of two neighbouring samples for each sensor as a function of time is shown in Fig. 15 for object sensors and in Fig. 17 for ego-motion sensors. The ground truth DGPS and consumer grade GPS are not included, since they report time-stamps with exact period 10 ms and 200 ms, respectively, so the graphs would show only a constant.

All sensors show stable period with expected values. There are fluctuations introduced by acquiring process and time-stamp quantization. Note that the 4-Layer LiDAR data frames are already fused by the acquisition system, therefore all four sensors share the same time-stamp.

The wheel odometry ego-motion sensor shows stable period with larger amount of dropouts. The period of consumer grade IMU is stable, with 1 ms quantization level.

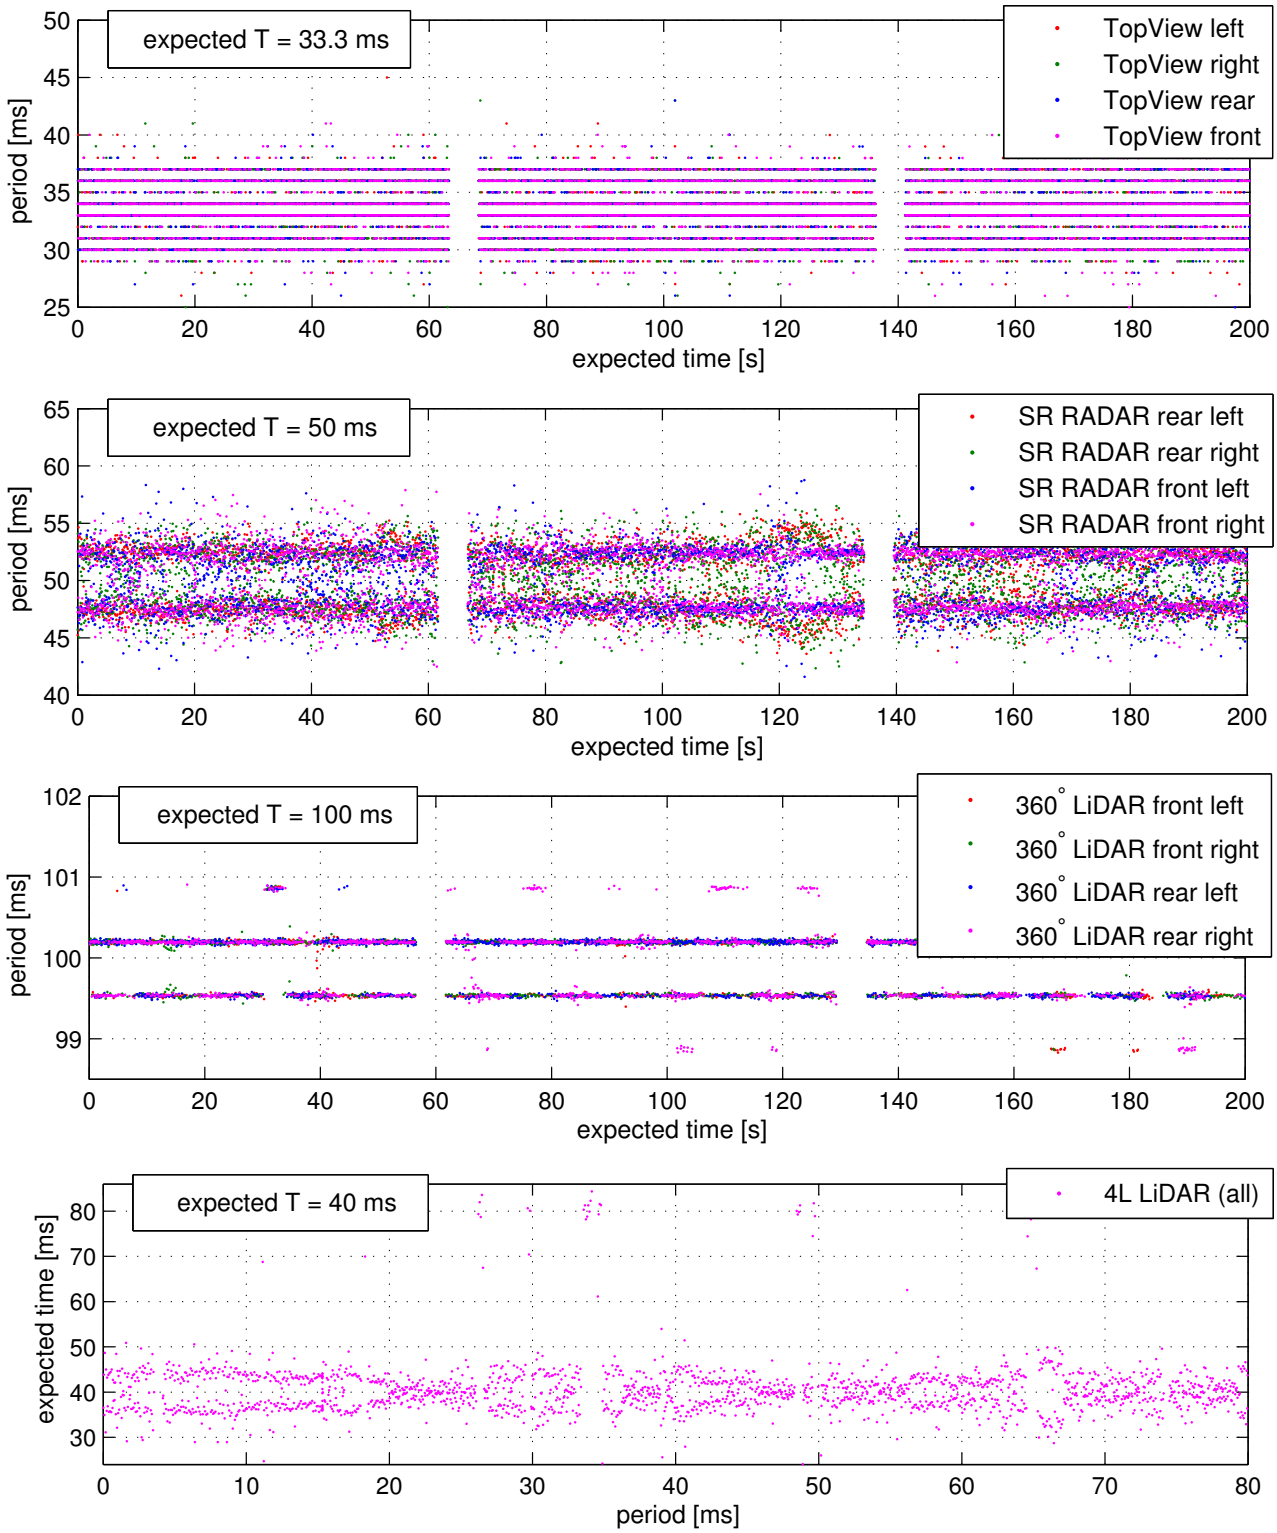

Figure 15: Period between arrival of two neighbouring samples for four principal types of sensors. Time-stamps for 360° LiDARs are hardware-based, time-stamps for the other sensors are obtained at the time of arrival to the data processor. Data from 4-Layer LiDAR comes from a different sequence than the others.

1.3.2 Expected time difference

To analyse long-term period stability and synchronicity the expected time-stamps is computed and compared with acquired time-stamps. Using the expected value of period, drop-

outs in data samples are detected from measured period between neighbouring data samples. This allows to reconstruct the index (order) number of each data sample. The expected correct time-stamp is the computed by multiplying the index with expected period and aligned with the beginning time of acquisition. The differences between measured and expected time-stamps are shown in Fig. 16 for object sensors and in Fig. 18 for ego-motion sensors.

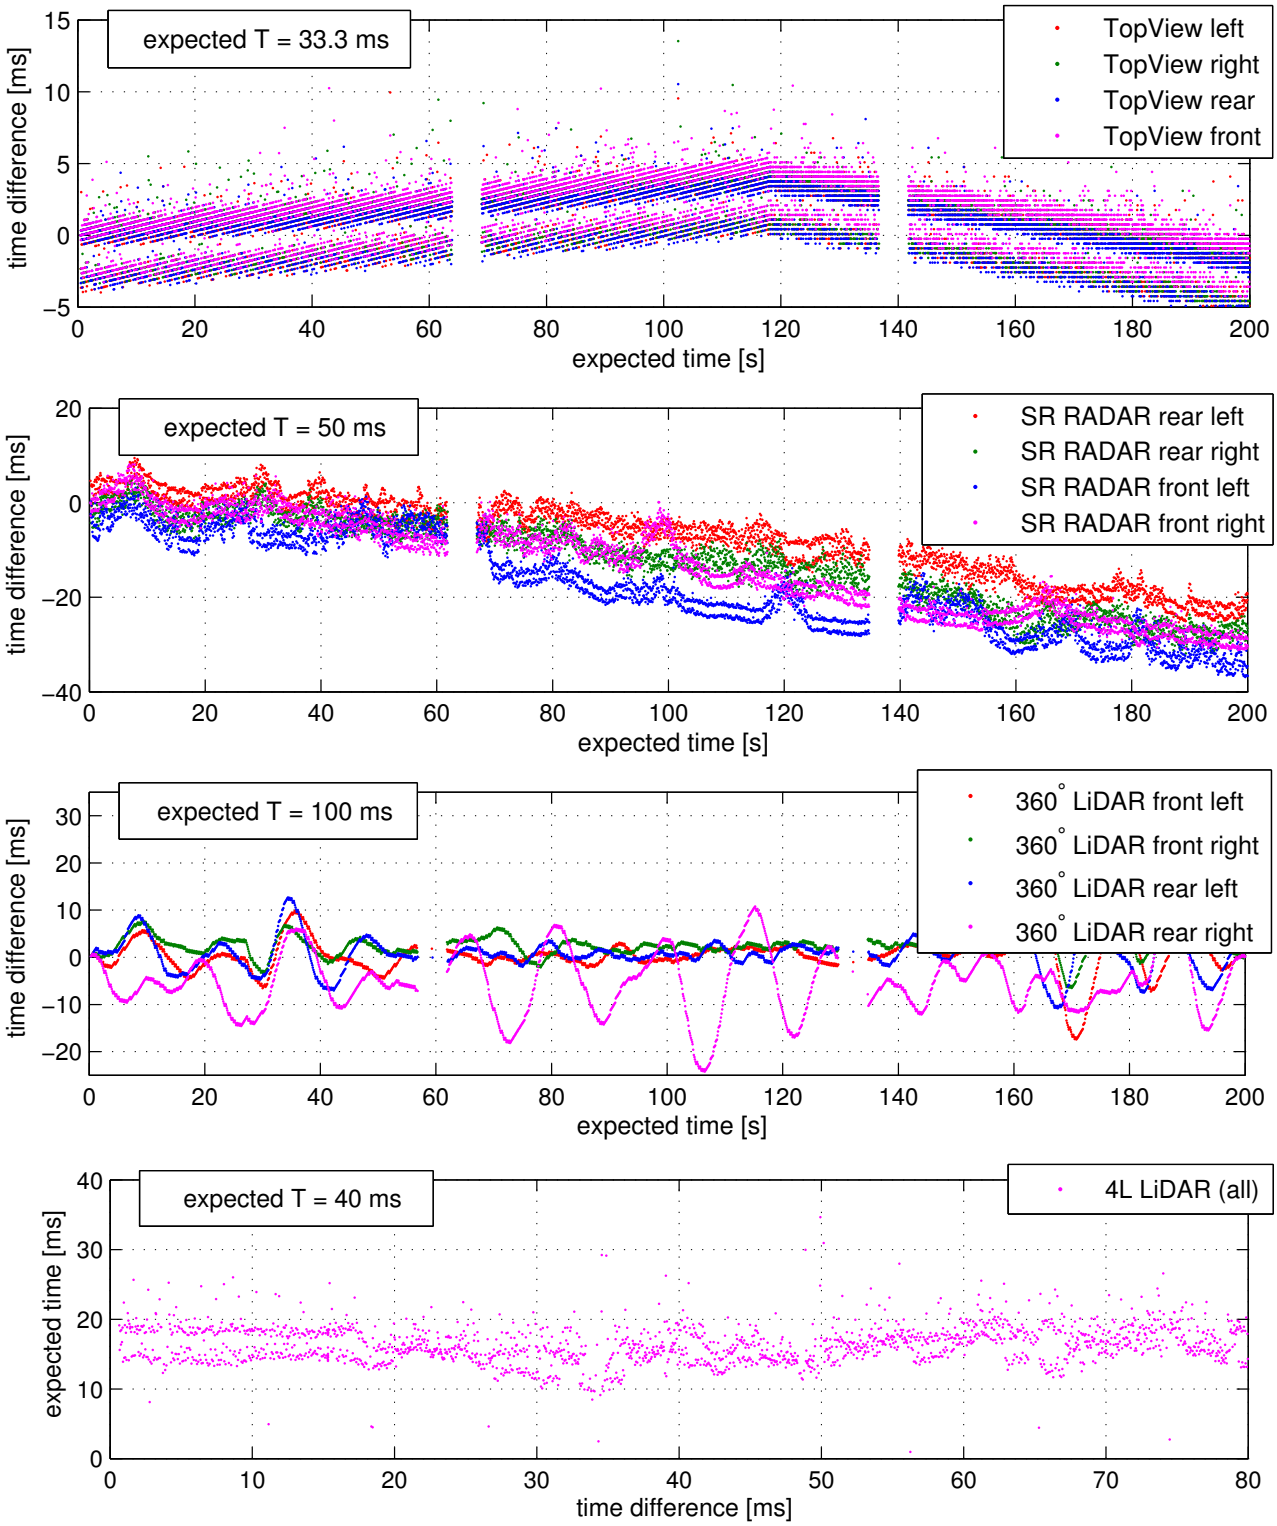

Figure 16: Stability of time-stamps. Difference between acquired and ideal data-sample time-stamp. The ideal time-stamp is derived from a data sample index and nominal period of a particular sensor. The data are for the same sequences as in Fig. 15.

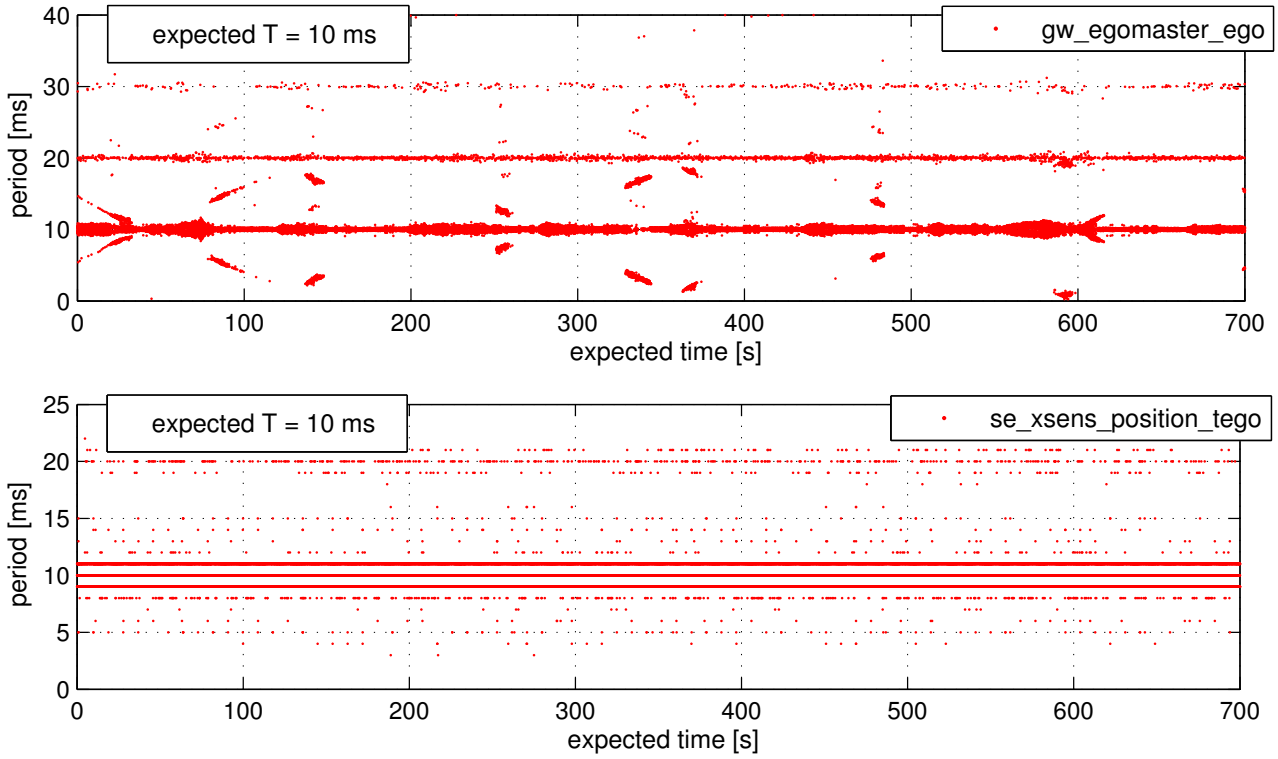

Figure 17: Period between arrival of two neighbouring samples for wheel odometry ego-sensor and Xsens ego-sensor. The two GPS based sensors report precise constant period without fluctuations and thus are not included here.

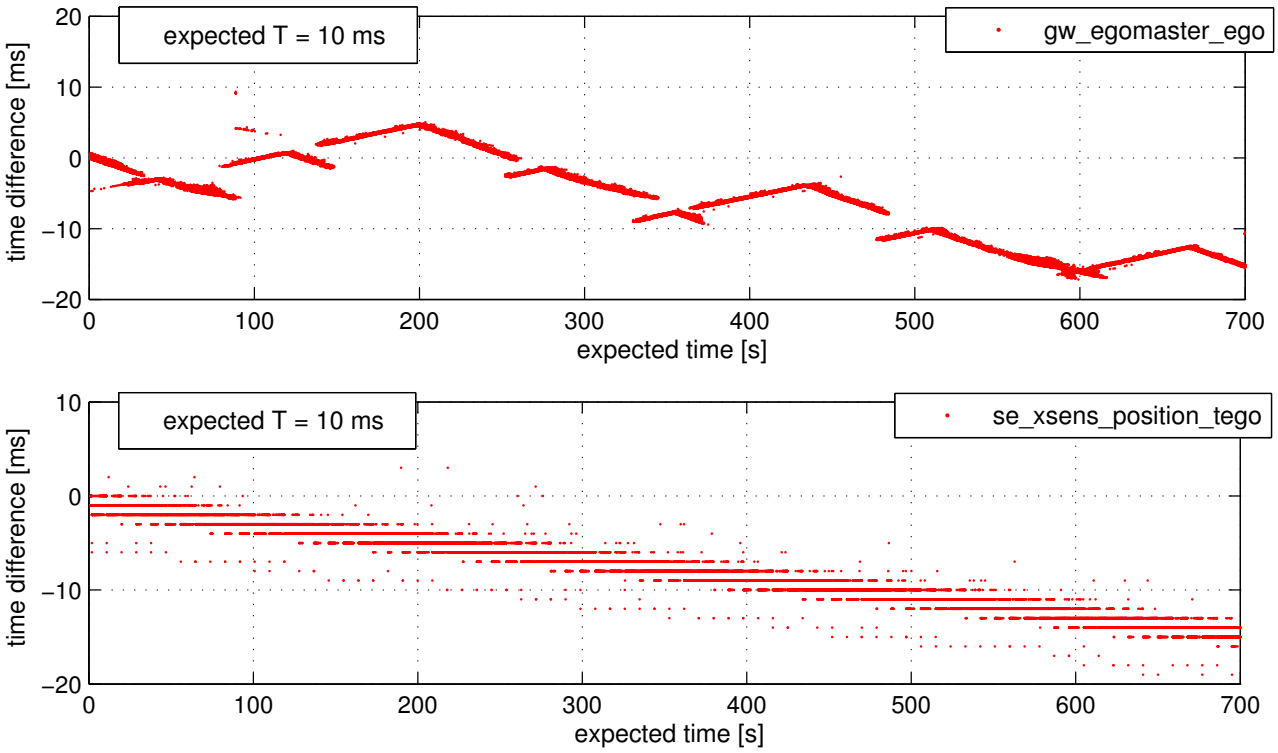

Figure 18: Stability of time-stamps. The figure has same meaning as Fig. 16. The two GPS based sensors report precise constant zero variation here and thus are not included.

We can see that the TopView cameras show some drift, corrected in a long term. The ethernet time synchronisation is assumed source. However, all cameras show the same tendency due to the fact that they are synchronised.

The 360° LiDARs show a fluctuations caused by rotor control loop, but the difference in long term keeps near zero, which means that the period is stable and the sensor are synchronous each other as well.

The differences of time-stamps of 4-Layer LiDARs are of the same order as in 360° LiDARs. They could be explained in a similar way, since there are rotating parts as well. The quasi-random noise may be caused by the fact that the data processing system needs to consider timestamps of all sensors but the sensors are not sychronized and occasional small slips occur. The system may also have an on-line adapting mechanism for selecting time instants that are used for motion correction.

The SR RADARs reveal some drift from zero and from each other a bit.

The wheel odometry ego-motion sensor shows a drift, periodically compensated. Both wheel odometry and consumer grade IMU sensors reveal a small long term drift 20 μs/s. This can be caused by different time-base of acquiring PC and must be still analysed.

2 Higher level processing framework

In this section we describe the tools developed with the goal of simplifying the software integration and error analysis in the car. Those tools include:

- data-recording system designed to always run in the background
- system health monitoring
- configuration and distribution of vehicle parameters

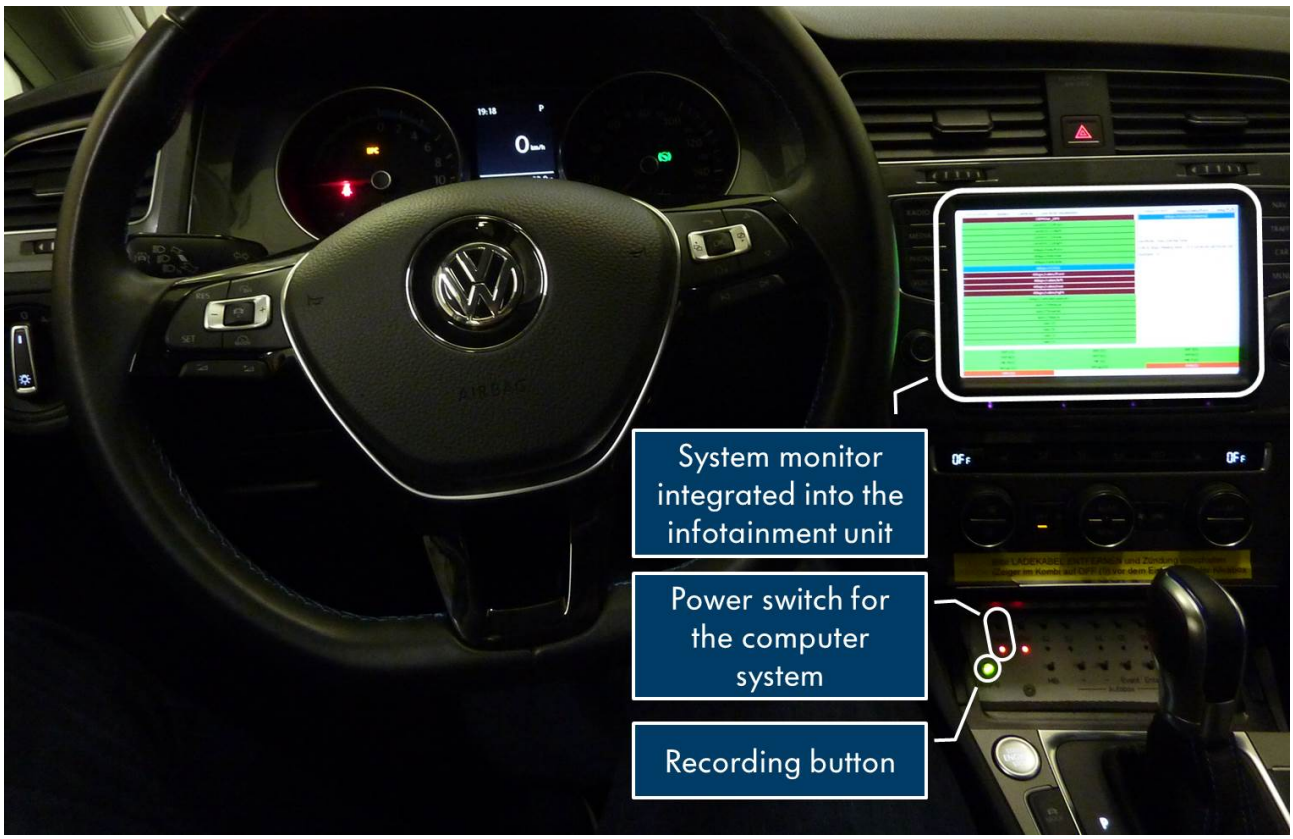

Figure 19: The most important elements integrated into the UP-Drive cockpit

Fig. 19 presents the cockpit of the UP-Drive car. Apart from the power switch for the computer system and the recording start/stop toggle button, also the system health monitor is shown. It is integrated into the infotainment unit. Thanks to those components, it is possible to perform some simpler tests or data-recordings without even needing to open the trunk of the car.

2.0.1 Data recorder

There are two use-cases for recording car data: planned data-set collections and spontaneous recording of data in case of relevant events such as system malfunctions or rare traffic situations. Both use-cases are handled using the same recording solution. It consists of data-loggers installed on 5 machines. One of them is dedicated to recording all of the relevant data - with the exception of image data. It is called 'central-logging'. 4 remaining

data-loggers are devoted to recording image data from the top-view system and each of the subsystems of the trifocal camera. The splitting of the recording to different machines is a design decision aiming at keeping the network traffic at an acceptable level. Of course, the data streams in all 5 of the recorders are synchronised to the same base time provided by an in-vehicle NTP server. The data-rates for the different data-loggers are summarized in Table 1.

Table 1: Data rates within the recording system

Recorder	MB/s	GB/h
central-logging	19,3	68
topview	10,4	37
trifocal-1	30,6	108
trifocal-2	26,6	94
trifocal-3	29,7	104
total	116,6	410

The recording can be triggered (and stopped) from within the drivers-seat by pressing the recording toggle key. Feedback is given to the driver in the form of an LED integrated into the recording key. In order to address the use case of recording spontaneous events, a ring buffer has been implemented and is set to 30 seconds. Currently all 5 recorders are triggered by the same key. Data can be merged into one file using the ADTF GUI tools or ADTF streaming library.

2.1 System health monitoring

For any complex system, its health monitoring becomes an important issue. In order to address it, we have developed the so called *system monitor* and enabled its use from the driver's or passanger's seat. Currently, the system monitor focuses mostly on displaying the health status of all the relevant hardware items in the car:

- computer system - including temperature, IP-connectivity and usage of ressources such as CPU, memory, network bandwidth or diskspace
- sensor system - including liveliness of sensors and the datarates or latencies of the provided data streams as well as visualisation of provided data (i.e. camera images)
- car system - including status of the vehicle ECUs or the power supply

As the monitored system is very complex, arranging information in an intuitive and readable way becomes a challenge of its own. The architecture of the system monitor allows for grouping information into tabs, and offers many different ways to display information. The 3 most relavant methods are:

- *traffic light display* - using different colors to signal the health status of a compontent
- providing text messages
- providing arbitrary graphic visualisations in form of video streams (camera images, 3D visualisation)

Figure 20 shows photos of sample system monitor screens integrated into the display of the navigation system of the car. In the future, additional views for all relevant software modules will be added.

Figure 20: Sample tabs of the status monitor: CPU health and usage, status of the sensor system, images from individual cameras

2.2 Configuration and distribution of vehicle parameters

As the UP-Drive sensor system consists of 23 environment sensors and 4 different systems providing ego-motion, it is very important to have unified mechanism of storing and accessing the calibration data. For this purpose, all of the calibration parameters - as well as other vehicle parameters (such as wheelbase, vehicle contour) - are arranged in a central XML file, which is versioned in a Git repository. In the vehicle, that XML file is always provided over a network share and is accessible by all the software modules running in the car. In this way all the modules have the possibility to use the same calibration parameters.

3 Communication capabilities

The computers in the trunk of the vehicle - as well as some sensors - are connected by Gigabit Ethernet networks. Most important of those is the network spanning the general purpose PCs. Its network load is presented in Fig. 21.

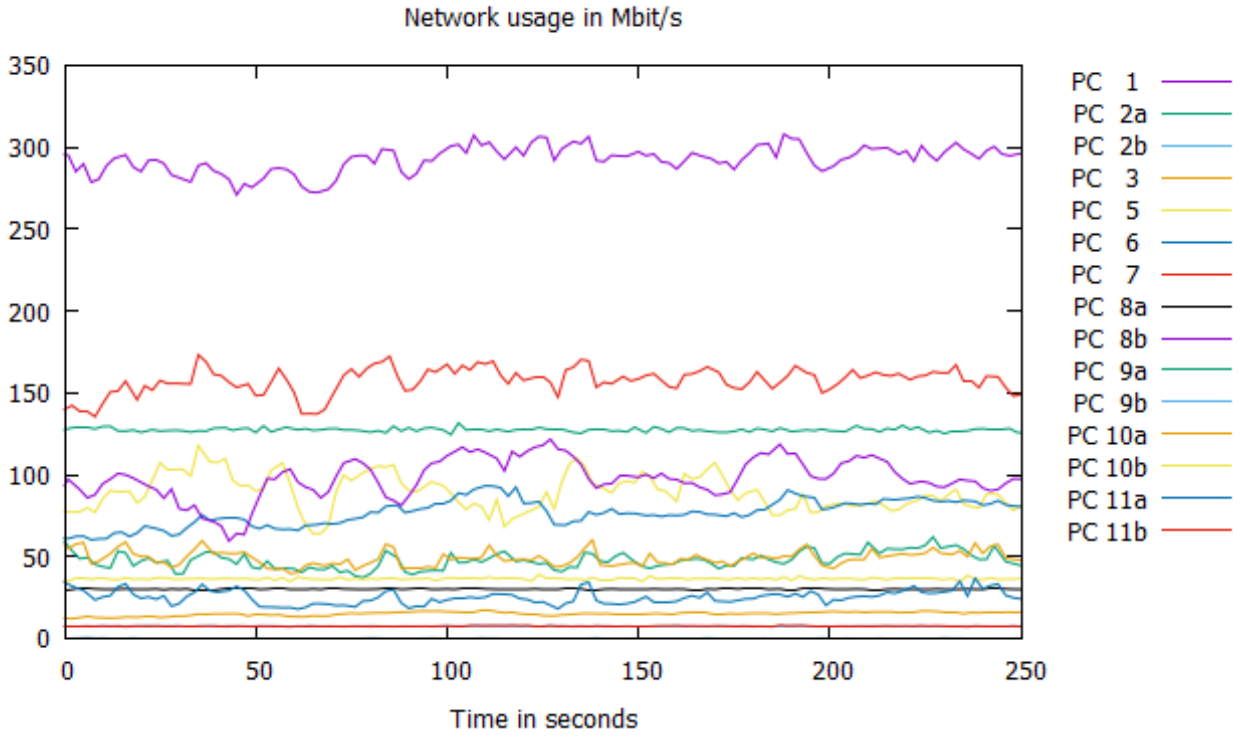

Figure 21: The network load on the PCs of the UP-Drive car

The presented data-rates have been captured for a software configuration used to log all sensor data. The highest network load is observed on PC 1, which is the PC used for logging the data. This is in agreement with our expectation. The data-rate stays at around 300 Mbit/s, which is below one third of the maximum theoretical bandwidth.

As the processed data (not yet included) is usually less in volume than the sensor data, the above result suggests that the network has enough bandwidth to accommodate for all the traffic to be generated in the system. However, since for stable and low-latency communication not only the average load but also the burst traffic is important, some care will need to be taken to design the traffic flow in an efficient way.

4 Safety elements and precautions

The car has been designed and built in a way that the safety driver can always take control of the car. Specifically, whenever the driver grabs the steering wheel or presses the acceleration or breaking pedal, the computer system gets disengaged from control of the vehicle.

In addition to that, the project has adopted the following safety policies:

1. All testing needs to be performed with safety driver in driver's seat.
2. Safety driver needs to have passed a safety training.
3. Safety driver needs to have been instructed in the specifics of the safety concept of the car.

In order to allow the possibility of doing demonstrations with no person in driver's seat, the passenger's seat has been equipped with braking and acceleration pedals. It yet needs to be evaluated though, whether the safety operator can control the car from the passenger's seat and what additional restrictions might be necessary.

5 Conclusions

This document shows that the first test vehicle is fully operational and satisfies all the defined requirements.

First of all, the sensor system has been extensively analyzed and found to work in accordance to specifications and the expectations. This includes the 3D sensors as well as ego-motion sensors. Data quality as well as integrity of the timing has been analyzed. It has also been shown that the sensors are correctly calibrated.

Second, the functionality of the high-level processing framework has been presented. Data recording as well as status monitoring are in place and can be expanded to accommodate the growing functionality of the system.

Third, the communication capabilities have been reported on. Finally, the safety concept was briefly explained.

Given the fact that the architecture of the first car has proven to fulfill the requirements, we plan to build the second car identical to the first by the end of the second project year.