Trade-Off between Coverage and Robustness of Automotive Environment Sensor Systems

Thomas Herpel, Christoph Lauer, Reinhard German
Computer Networks and Communication Systems
University of Erlangen-Nuernberg
91058 Erlangen, Germany
{herpel, german}@informatik.uni-erlangen.de
christoph.lauer@informatik.uni-erlangen.de

Johannes Salzberger
Department of Safety Electronics
Audi AG
85045 Ingolstadt, Germany
johannes.salzberger@audi.de

***Abstract*—Car manufacturers, research institutions and governments are challenged by decreasing the number of road traffic fatalities in spite of rapidly increasing traffic volumes. One approach is to enhance the "intelligence" of modern cars in anticipating dangerous driving situations. Besides wireless communications with the traffic environment, context perception enabled by well-known technologies such as radar, laser or video is regarded as a promising instrument in driver assistance and collision avoidance. There is a rich variety of both hardware configurations of automotive environment sensor systems and strategies for sensor data fusion and processing. For sensor systems assembled of multiple devices, the trade-off between increased spatial coverage, robustness against interfering factors and accuracy of measured data is a crucial aspect of system design. We investigate up-to-date automotive sensor systems with respect to detection performance and detection quality. We built UML-based discrete-event simulation models to resemble sensors, context perception, signal processing and data fusion in realistic road traffic scenarios. Our results allow for a sound comparison on sensor systems and data acquisition strategies at an early stage of system development.**

Keywords-Environment Sensors, Sensor Data Fusion, Simulation, Driver Assistance, Automotive Applications

I. INTRODUCTION

In recent years, industrialized countries experienced a significant increase in traffic volume. Besides on serious issues of environmental protection – reducing CO2-emissions is the buzzword – lots of effort is spent on concepts for enhanced road safety. The major challenge for car manufacturers, research institutions and governments worldwide is "to keep the cars from crashing" and thus to avoid fatal car accidents or mitigate the crash severity [1]. In addition to concepts with a global view on transportation systems like car-to-car communications, intelligent roadside infrastructures or advanced traffic routing and information services, enhancing the "intelligence" of individual vehicles within the traffic flow is regarded as a key-idea for proactive safety systems. To achieve this, sensor technologies well-known from other application areas like military or civil aviation found the way into the repertory of on-board electronics. Radar, laser, ultrasonic or video sensors perceive information about the environment and possible threats around the vehicle, anticipate dangerous driving situations and activate adequate protection means if necessary. The set of environment sensors, electronic control units, algorithmic functions and actuators is often referred to as Advanced Driver Assistance Systems (ADAS). Figure 1 shows various ADAS applications in a possible future road traffic scenario.

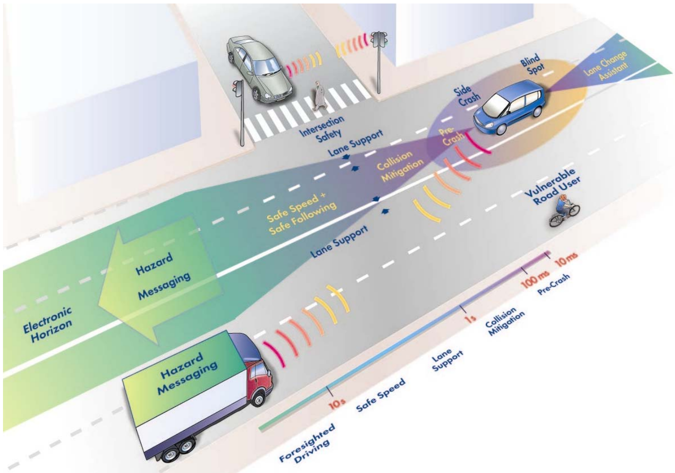

Figure 1: ADAS and Road User Communication in Future Traffic Scenario[1]

As single sensor systems are likely to exhibit undesired weaknesses, a sensor set for safety critical ADAS applications typically consists of multiple devices. The individual sensors may differ in measuring principle, detection range, field-of-view (FOV), weather robustness or other attributes. Hence, sensor signal processing and data fusion of multiple devices is a sophisticated process including important design decisions regarding system performance and dependability. First and foremost, there are different levels of abstraction on which data from various sensors is consolidated. Algorithms may be applied to either near-raw sensor data at a very early stage of signal processing (low-level data fusion) or on individually pre-processed sensor data represented as object lists (high-level data fusion) [2]. Both approaches are expected to have certain assets and drawbacks in terms of entropy of information and computational complexity. In addition, on both levels the sensor data can be combined focussing on increased coverage

[1] Taken from http://www.prevent-ip.org/

or on increased confidence. In the former approach, data from sensors with partly disjoint FOV enlarges the overall coverage area, whereas in the latter procedure data from sensors with partly joint FOV is used to validate object detections of the other sensor(s). Various investigations on sensor hardware and sensor data processing have been carried out, for example on radar image acquisition [3], multi-target tracking [4], sensor-based cruise control [5], pedestrian protection [6] and devices for ADAS functions [7]. However, research is mainly performed as cost intensive and manpower consuming road tests. The benefits of modelling and simulation of environment sensor systems are up to now exploited rarely [8]. We use discrete-event simulation and real-time Unified Modeling Language (UML) to model environment sensor systems as well as low- and high-level data fusion and strategies of increased coverage and increased confidence in realistic traffic scenarios. The simulation results allow for a comparison on what data combination strategy performs best in which scenario and is preferable with respect to maximum detection performance and dependability of ADAS applications. The paper is organised as follows: Chapter 2 introduces to ADAS architectures and chapter 3 describes the associated simulation model. The simulation results are presented in chapter 4. Finally, chapter 5 concludes this paper and presents some areas of future work.

II. ADAS ARCHITECTURES

The relevant sensor technologies for environment perception, techniques for signal processing and common approaches for sensor data fusion are presented in the following.

A. Environment Sensor Systems

Sensor devices for context perception are well proven from environment surveillance in military or civil aviation. In the recent past, some of these technologies were adopted for automotive applications of proactive systems, i.e. systems, which rely on environment information from sensors and are able to (partly autonomous) assist the driver in complex driving situations. Depending on how information is perceived, the employed devices can be categorized into active sensors and passive sensors.

As the name implies, active sensors are actively probing the environment by emission of beams of certain wavelength. Object detection depends on the reflection of the emitted beams at the object's surface. The sensor directly derives the distance to the object by measuring the time-of-flight between pulse emission and pulse response. Commonly known active sensor technologies are radar, laser or ultra-sonic. Radar devices in automotive applications can be divided into short range radars (SRR) and long range radars (LRR). Typical frequencies, ranges and FOV values are 24 GHz, 20-50m and 20-60° for SRR and 76 GHz, 100-150m and 5-8° for LRR devices. In addition to an excellent robustness against bad weather (rain, fog, snow or spindrift) and an accurate longitudinal distance sensing, radar has the intrinsic capability to determine an object's velocity by evaluating the Doppler shift in frequency between emitted and reflected signal. A drawback of radar sensors is that only very few reflection points along edges are generated, which complicates assessing the lateral extension of an object and makes radar prone to misinterpretations, especially in scenarios with high traffic density. Laser based devices, e.g. lidar sensors, laser-scanners (LS), photonic mixer devices (PMD) or closing velocity sensors (CV), show typical ranges from about 20m (CV) up to 120m (lidar, LS) with FOV values from 10° (lidar) up to 360° (LS). These sensors are more accurate in lateral object detection than radar devices, as significantly more reflection points along an object's edges are generated. Thus, pattern matching approaches can be applied to determine the exact object geometry and to figure out clutter objects, i.e. objects like trash cans or manhole covers, which are physically present but nevertheless uninteresting to an ADAS application. Disadvantages of laser sensors are the susceptibility to bad weather and the missing way of direct measuring of an object's velocity, which has to be derived by tracking the movement along time instead. Acoustic devices like ultra-sonic sensors have been widely used in automotive applications like park distance controls. Due to the very short sensing ranges and negative influences of loud driving noise, these sensors play a minor role in safety critical ADAS.

In contrast to active sensors, passive devices do not emit any probing signal but passively perceive the environment comparable to the scene perception by the driver's eyes. Especially vision-based systems like mono-, stereo- or infra-red-cameras are currently investigated for applications in traffic sign recognition or night-vision. The sensors offer excellent means for target classification as pattern matching algorithms can easily be applied to the pixel-based scene representation. A drawback of vision sensors is the inaccurate distance information and the sensitivity against dirty weather, especially due to the installation location between inside rear view mirror and front windshield. For most current ADAS applications, vision-based sensors are used to validate detections from active sensors by searching in a so-called region-of-interest within the pixels. In the following, we will focus on radar- and laser-based active sensors and automotive environment sensor systems composed of such devices.

B. System Architecture

Sensor data from multiple devices passes a sequence of steps in data fusion and signal processing (cf. figure 2). First of all, all sensors contributing data to the ADAS need to be synchronized in terms of a common temporal and spatial base. In a low-level fusion, the synchronized near-raw sensor data is subsequently fused on a low level of abstraction before the signal processing algorithms are applied to the conglomerate of measured values. By contrast, in high-level fusion architectures each sensor individually pre-processes data and the actual data fusion is then carried out on a higher level of abstraction, e.g. on object lists [9].

Obviously, a high-level fusion allows for individually adapted signal processing algorithms whereas for a low-level fusion more generic settings have to be found in order to process heterogeneous sensor data along a single path. In the clustering step, groups of data points are assembled that are believed to be originating from the same real world entity. Feature extraction derives relevant attributes of the clustered entities like distance, velocity, acceleration and width. In the association step, a record of objects from the last time step correlated with the currently present objects is established. The filter algorithm lessens the effect of sensor inaccuracy and signal noise and provides the state predictions for the next iteration's association step. Via tracking, the set of objects is updated, which means new objects are added, vanished objects are removed or states of present objects are adjusted.

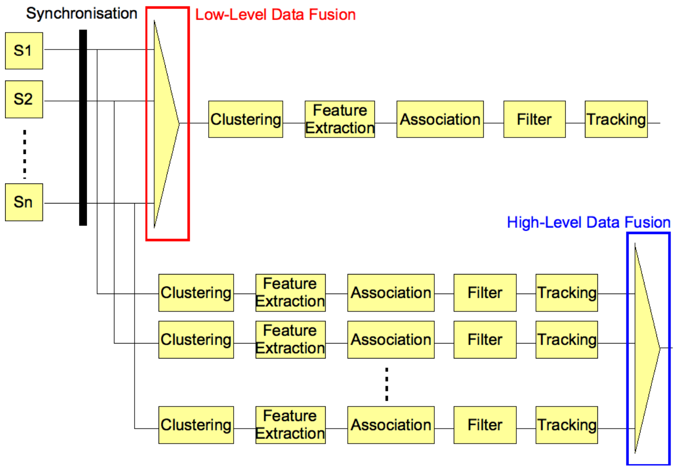

Figure 2: Sensor Signal Processing with Low- and High-Level Data Fusion

C. *Signal Processing Algorithms*

Various algorithms have been proposed for signal processing. Clustering mostly works hierarchical, either divisive, i.e. the entire set of data points is considered as one cluster and separated into subsets, or agglomerative, starting with each entity as a single cluster, building super-clusters by grouping similar entities together. Feature extraction transforms a set of measurements into a vector of features, for example by averaging over corresponding values within one cluster [10]. Algorithms for association of objects from different time steps include greedy approaches like Nearest Neighbour (NN), Joint Probabilistic Data Association (JPDA) or Multiple-Hypothesis Testing (MHT), varying in terms of computational load, local or global optimal solutions and the maximum number of objects per track association. The de-facto standard for filtering applications is the Kalman-filter [11]. In an iterative process the filter calculates a weighted update of the state prediction from the last iteration, based on the actual measurement and the Kalman gain matrix. The Kalman gain is derived from the signal covariance values. Other emerging filter approaches include, for example, the particle filter or interactive multi-model filters.

D. *Sensor Data Fusion Paradigms*

Sensor data from multiple sources can be fused in respect of different fusion paradigms on different levels of abstraction. Low-level fusion is performed early at a stage of near-raw data, whereas high-level fusion is carried out after individual pre-processing of data to a higher level of abstraction. Low-level fusion is assumed to be optimal in terms of information content of data to be fused [12], yet with the drawback of higher computational and communicational load and reduced modularity and scalability if large amounts of near-raw data have to be evaluated and communicated. Fusing data at higher levels increases the modularity at the interfaces, e.g. standardized object lists from different vendors of sensor devices, and allows for a more adapted processing of heterogeneous sensor data in the parallel branches. However, reduced complexity might come at the expense of reduced content of valuable information at higher levels. Not all sensor combinations can be fused on either low- or high-level. For example, a low-level fusion of pixel-based video data and reflection-based radar data is a fairly unprofitable process due to the significantly different characteristics of the measured values. Recent publications provided insight into fusion architectures, paradigms and case studies [3-7, 9, 12].

E. *Benefits of Sensor Data Fusion*

Sensor data fusion is expected to generate information of "greater quality" out of individual data from single sources [13]. For ADAS applications, there are two major benefits data from multiple sensors can be fused for. *Increased coverage* is achieved if data from sensors with (partially) disjoint FOV is fused in order to enlarge the overall coverage area. This so-called complementary fusion is comparable to a 'logical-OR' operation on sensor data and is often necessary to enable sufficient coverage for long and short range applications. The counterpart strategy yields *increased confidence* in environment information. Data from sensors with (partially) joint FOV is used to validate object detections of the other sensor(s). This approach is also called competitive fusion, comparable to a 'logical-AND' operation on individual observations.

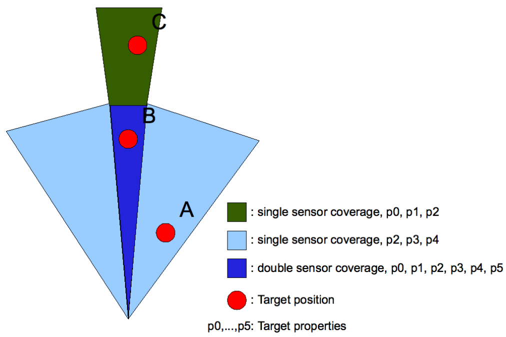

Figure 3: Coverage Strategies for Object Detection and Sensor Data Fusion

To clarify the difference, figure 3 shows a sensor set consisting of two sensors with partially overlapping coverage areas. The red dots resemble positions of a target with the abstract properties p0-p5, where p0, p1 and p2 are available to the system if the target is detected by the long-range sensor (green area), p2, p3 and p4 are provided by the short-range sensor (light blue area) and p5 is only available in the area of overlapping coverage (dark blue area). For a complementary fusion, it is sufficient that one of the sensors detects the target at any of the positions A, B and C. Thus, object detection is possible anywhere within both FOV, yet with a varying number of obtainable properties. In contrary, a competitive fusion restricts object detection to the overlapping area, which reduces the number of detectable positions of an object to position B, but in turn increases the trustworthiness and offers the maximum number of object properties. Safety critical ADAS applications require both high availability of the sensor system (in terms of minimum number of missed objects) and high dependability (by means of minimum number of false detections). Which ADAS architecture performs best in this case is investigated in the following.

III. Simulation of Environment Sensor Systems

Research on automotive environment sensor systems mostly relies on experiments with dedicated hardware installed in test cars. As authentic and valuable this approach is it lacks flexibility in on-the-fly replacement of fusion paradigms and combinations of environment sensors. In this paper, we use discrete-event simulation to evaluate data fusion architectures [14]. The simulation model was created using the multi-method simulation tool AnyLogic 5.5.4 [2], enabling real-time UML behavioral descriptions and full Java support.

A. Simulation Model

The objective of the simulation model is to resemble closely the central components of an automotive environment sensor system. This incorporates sensor hardware, signal processing algorithms and a realistic road traffic environment, which is capable of generating appropriate stimuli for the sensor system. The UML diagram in figure 4 shows the basic model structure and the data dependencies.

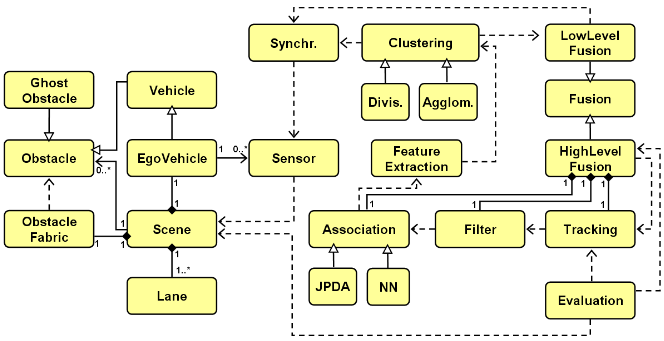

Figure 4: UML Diagram of Environment Sensor System Simulation Model

The central components are the scene generator, the signal processing blocks and the fusion and evaluation classes. The scene generator is responsible for spawning relevant and irrelevant objects into the context of the ego-vehicle – this is the vehicle where the sensor system is installed – and creating reflection points, taking into account the individual sensor properties, namely FOV, range, lateral resolution, sensor variance and measuring principle. Signal processing and data fusion elements perform environment perception based on these reflection points. High-level fusion and low-level fusion are evaluated simultaneously, considering ground truth information as provided by the scene generator. The modular structure of the simulation model permits experiments with various combinations of sensor devices, traffic scenarios and fusion paradigms and provides excellent means for analyses of automotive environment sensor systems.

[2] http://www.xjtek.com/anylogic/

B. Configuration of Simulation Experiments

We conducted several simulation experiments and evaluated the performance of various environment sensor systems. Table 1 shows the scenario settings for urban and motorway traffic. Spawn times for objects are sampled from a uniform distribution. A specific fraction of objects is generated as clutter objects, which are irrelevant to the system.

TABLE I. Parameter Settings for Simulated Traffic Scenarios

	Urban	**Motorway**
Object spawn time (s)	uniform (0, 0.3)	uniform (0, 0.12)
Fraction of clutter objects	0.45	0.35
Max. number of obj. in scene	5	6
Number of lanes	2	3
Lane width (m)	3	3.5
Scene horizon (m)	50	80
Ego-vehicle speed (km/h)	40	110
Lane speed difference (km/h)	10	20

TABLE II. Parameter Settings for Sensors in Sensor Sets

Sensor	**Range (m)**	**Azimuth**	σ_x	σ_y	σ_{vrel}
LRR (set 1)	120	9°	0.45	0.15	0.3
LRR (set 3, 4)	120	10°	0.45	0.15	0.3
LS (set 1)	100	100°	0.11	0.11	--
LS (set 5)	28	120°	0.11	0.11	--
Lidar (set 2)	120	10°	0.15	0.15	--
SRR (set 2)	40	60°	0.5	0.2	0.5
SRR (set 7, 8)	28	60°	0.5	0.2	0.5
PMD (set 3)	40	55°	0.25	0.25	--
PMD (set 6, 8)	28	55°	0.25	0.25	--
CV (set 4)	20	45°	--	0.15	--
CV (set 5, 6, 7)	12	24°	--	0.15	--

Table 2 shows the parameterization of the sensor devices in the respective sensor sets. "Range" denotes the upper limit of longitudinal detection ability and "Azimuth" the FOV. The σ-values refer to the setting of the Gaussian probability distribution with zero mean, modeling the sensor variance in longitudinal (σ_x), lateral (σ_y) and velocity (σ_{vrel}) detection. The higher the value, the more inaccurate is the sensing. We arranged eight environment sensor systems as denoted in table 3. Sets 1-4 are combined long- and short-range systems whereas sets 5-8 focus on close-up range surveillance.

TABLE III. Configurations of Simulated Sensor Sets

	Long-range device(s)	**Short-range device(s)**
Sensor set 1	2xLRR[3]	LS
Sensor set 2	Lidar	SRR
Sensor set 3	LRR	PMD
Sensor set 4	LRR	CV
Sensor set 5	--	LS, CV
Sensor set 6	--	PMD, CV
Sensor set 7	--	SRR, CV
Sensor set 8	--	SRR, PMD

[3] Installed left and right in bumper with a horizontal offset

Sensor data is processed and fused according to the strategies described in chapter 2, using agglomerative clustering, generic feature extraction, JPDA and Kalman filtering. Both low-level and high-level data fusion are carried out applying the increased coverage and the increased confidence strategy.

IV. RESULTS

The simulation results were obtained by at least ten independent replications of a single experiment, applying a simulation control with a 95% confidence interval and at most 10% of relative error in observed performance measures between consecutive simulations [14].

To determine the *accuracy* of the sensor system, we evaluated the mean overall error in detected object position (in meters). It is computed as the mean value of all Euclidian Distances between actual coordinates of objects in the scene and the corresponding positioning information provided by the multi-sensor system. The system's *availability* is measured by the mean false negative detection ratio. "False Negative" means that the sensors oversaw an object, which the system was expected to detect with respect to actual settings for hardware and algorithms. For example, if 4 objects are present in the scene and only 3 of them are detected, the ratio is 25%. Of course, the higher the mean false negative detection ratio, the fewer is the information about the environment. The *dependability* of the sensor system is expressed by the mean false positive detection ratio. It resembles the robustness of the system against misinterpretations of the scene and thus propagation of irrelevant objects or creation of ghost objects, i.e. objects which are actually not physically present in the scene. For instance, a mean false positive detection ratio of 20% is obtained if the sensor system assumes 5 objects to be in the environment, although only 4 objects are really present. Thus, 20% of the propagated objects are virtually created, e.g. due to mismatching in clustering or association. For greater clearness, figure 5 depicts a constellation of environment sensing with sensors S1 and S2 with partly overlapping FOV.

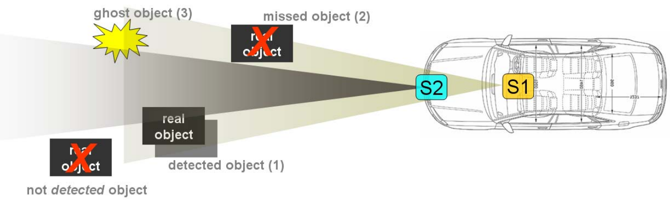

Figure 5: Accuracy (1), Availability (2) and Dependability (3) of Sensors

The outcomes for the mean detection error for increased coverage and increased confidence are shown in figures 6 and 7. The figures are subdivided into low-level fusion (left hand side) and high-level fusion (right hand side), with colored bars depicting the results for the eight individual sensors sets (cf. table 2, 3 and captions). Obviously, the error is consequently higher for increased coverage compared to increased confidence, where propagated object information is validated by multiple sensors. In addition, for both strategies low-level fusion mostly provides more accurate positioning information than high-level fusion, indicating a higher content of information about objects at a lower level of abstraction.

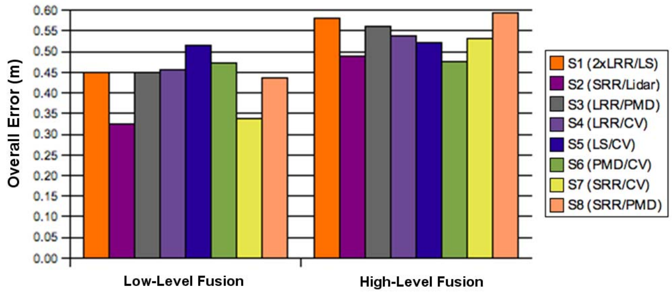

Figure 6: Mean Overall Error in Detection Accuracy (Increased Coverage)

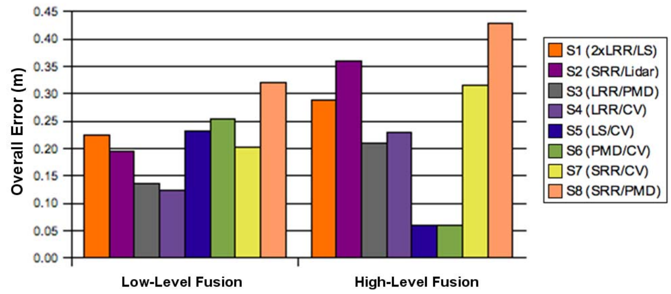

Figure 7: Mean Overall Error in Detection Accuracy (Increased Confidence)

In figures 8 and 9, the results for mean false negative and mean false positive detection ratio are depicted, with low-level fusion results in the two left columns and the corresponding values for high-level fusion in the right columns. The colored bars represent the eight different sensor sets. Increased coverage yields a considerably reduced number of false negative detections compared to increased confidence. In the latter strategy, the requirement of object detection by all sensors and mutual validation leads to more missed out targets, detected by only one sensor device.

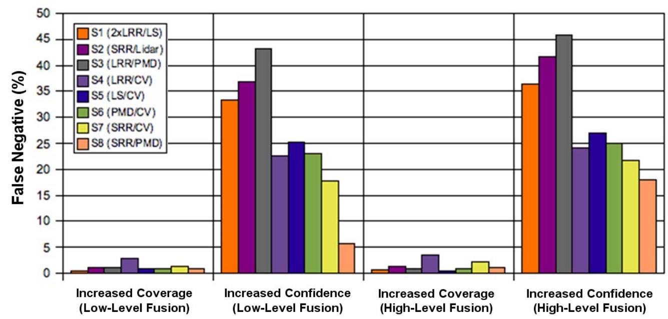

Figure 8: Mean False Negative Detection Ratio (Availability)

The direct opposite can be observed for the characteristics of the mean false positive detections as shown in figure 9. The increased robustness strategy filters out significantly more clutter objects and ghost objects than the strategy of increased coverage. Again, low-level fusion performs in most cases better than fusing the data at higher levels of abstraction. The two-way confirmation of detection is the central point in avoidance of false alarms. For instance, a radar sensor detects a reflecting metallic beverage can on the street as a relevant object, whereas a laser-based sensor would sense the lacking vertical extension, classify the target as irrelevant and outvote the radar in a 'logical-AND' combination of the observations.

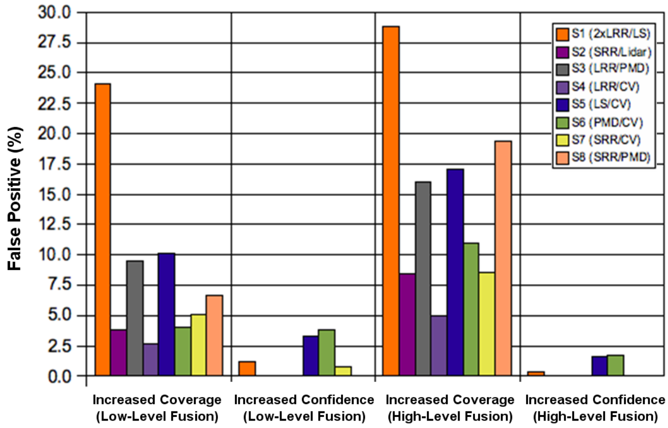

Figure 9: Mean False Positive Detection Ratio (Dependability)

In all configurations, the deviations between the simulation results resemble fairly well the peculiarities of the respective systems. For example, sensor set 1 with two long-range radar lobes and a laser-scanner features an extensive FOV, thus, at an increased coverage strategy it is more prone to false positive detections in both low-level and high-level data fusion and in turn less exposed to false negative detections. Sensor set 8, consisting of heterogeneous technologies (SRR and laser-based PMD), performs exceptionally well at the increased robustness strategy in terms of false negative and false positive detections (cf. figures 8, 9).

V. Conclusions and Future Work

We investigated automotive environment sensor systems for ADAS applications with respect to accuracy in object detection, availability and dependability of the overall system. We varied the data fusion strategy on both the level of abstraction (near-raw data at low level, object lists at high level) and the logical level of sensor data combination (increased coverage, increased confidence). We used UML-based discrete-event simulation and built a realistic model for sensor devices, signal processing, data fusion and road traffic scenarios. The simulative approach and the modularity of the model allowed us to conduct several meaningful experiments to figure out strengths and weaknesses of various system configurations. We analyzed the simulation output thoroughly and we can draw the conclusion that a competitive fusion of sensor data which was obtained by different measuring principles has a considerably positive impact on accuracy and dependability. However, the availability decreases at the same time, as detection of objects is essentially restricted to the overlapping area in FOV. Investigating strategies for an early, low-level data fusion appears to be worthwhile, as this approach performs better than high-level fusion in most constellations. The sensor systems we investigated so far revealed no cure-all for the crux in the trade-off between detecting everything with the maximum possible accuracy, yet not detecting anything irrelevant, though. Hence, a direction for future work on this topic is to include other technologies in the simulation, for instance passive vision-based sensors or data dissemination techniques in car-to-car scenarios. Engineers from Audi attested our observations and outcomes a close correlation to experiences in real hardware tests, thus it is desirable to work out test strategies with enhanced ground truth reference measurements for experimental validation of the simulation results.

References

[1] W.D. Jones, *Keeping cars from Crashing*, Spectrum, IEEE, vol. 38, no. 9, pp. 40-45, 2001.

[2] D.L. Hall and J. Llinas, *Handbook of multisensor data fusion*, CRC Press, Boca Raton, USA, 2001.

[3] U. Meis and R. Schneider, "Radar Image Acquisition and Interpretation for Automotive Applications," *Proc. of IEEE Intelligent Vehicle Symposium*, Columbus, USA, pp. 328-332, 2003.

[4] R. Mobus and B. Kolbe, "Multi-Target Multi-Object Tracking, Sensor Fusion of Radar and Infrared," *Proc. of IEEE Intelligent Vehicle Symposium*, Parma, Italy, pp. 732-737, 2004.

[5] G. Widmann et al., "Comparison of Lidar-Based and Radar-Based Adaptive Cruise Control Systems,“ *Proc. of SAE World Congress*, Detroit, USA, pp. 1-14, 2001.

[6] M. Meinecke et al., "Approach for Protection of Vulnerable Road Users Using Sensor Fusion Techniques,“ *Proc. of International Radar Symposium*, Dresden, Germany, 2003.

[7] J. Langheim et al., "CARSENSE – New Environment Sensing for Advanced Driver Assistance Systems,“ *Proc. of IEEE Intelligent Vehicle Symposium*, Tokyo, Japan, pp. 89-94, 2001.

[8] M. Bühren and B. Yang, "Simulation of Automotive Radar Target Lists Considering Clutter and Limited Resolution," *Proc. of International Radar Symposium*, Cologne, Germany, pp. 195-200, 2007.

[9] A. Amditis, *PReVENT Fusion Forum e-Journal*, vol. 2, 2008.

[10] N. Kämpchen et al., "Ein Sensorfusionssystem für automotive Sicherheits- und Komfortapplikationen," *Aktive Sicherheit durch Fahrassistenz*, 2004

[11] A. Gelb, *Applied Optimal Estimation*, MIT Press, 1974.

[12] T. Tatschke, "Early Fusion," *PReVENT Fusion Forum e-Journal*, vol. 1, pp. 6–7, 2006.

[13] L. Wald, "Definitions and terms of reference in data fusion," *Int. Archives of Photogrammetry and Remote Sensing*, vol. 32, no. 7, pp. 651–654, 1999.

[14] A.M. Law and W.D. Kelton, *Simulation Modeling and Analysis*, McGraw-Hill Higher Education, 1997.