CHALMERS
UNIVERSITY OF TECHNOLOGY

Sensor fusion and sensor management for automated vehicles

A way to win the Grand Cooperative Driving Challange

Master's thesis in Complex Adaptive Systems

MATS SVENSSON

Department of Mechanics and Maritime Sciences
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2019

Master's thesis in Complex Adaptive Systems

Sensor fusion and sensor management for automated vehicles

A way to win the Grand Cooperative Driving Challange

MATS SVENSSON

Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Automated Systems
Chalmers University of Technology
Gothenburg, Sweden 2019

Sensor fusion and sensor management for automated vehicles
A way to win the Grand Cooperative Driving Challenge
MATS SVENSSON

Master's Thesis 2019:91
Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Automated Systems
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: The truck from Chalmers Truck Team participating in the Grand Cooperative Driving Challenge 2016.

Printed by Chalmers Digitaltryck
Gothenburg, Sweden 2019

Sensor fusion and sensor management for automated vehicles
A way to win the Grand Cooperative Driving Challenge
MATS SVENSSON
Department of Mechanics and Maritime Sciences
Chalmers University of Technology

Abstract

The development of the autonomous vehicle is one of modern technology's widespread goals and is therefore an area of great competition and rapid development. This puts great pressure on technology as a whole, but on sensors and sensor fusion especially, since they are instrumental for the implementation of active safety. This thesis examines how sensors and sensor fusion is used in the different autonomous vehicles that participated in the 2016 Grand cooperative driving challenge (GCDC), and how it was implemented in one of them specifically. The conclusion is that the GCDC is a great incentive for the development of cooperative autonomous vehicles in general, although it lacks in the specific area of sensors and sensor fusion. However, the hardware and concepts used is an indication of what is needed for a vehicle to be fully equipped to deal with the rigors of the real world and to guarantee everyone's safety in traffic.

Keywords: Sensor fusion, 2016 Grand Cooperative Driving Challenge, Autonomous Vehicles, Kalman Filter

Acknowledgements

To begin with I want to thank two of my friends: Milica Bijelovic for convincing me to take on this, my second Masters Thesis and Emma Wijkmark, without whom this report would never have been finished.
A great thanks goes out to Ola Benderius, my supervisor during this thesis, for all support he has given me and the patience he has shown me. I also want to thank my colleague thesis workers Johan Bergström, Ivo Batkovic, Björn Persson Mattsson and Björnborg Nguyen for making the thesis time a tolerable and even fun experience.
In the end I want to thank my family and all my friends for making me who I am today and for supporting me during my, by now very long, time at Chalmers.

Mats Svensson, Gothenburg, June 2019

Contents

1
Introduction

The work presented in this thesis is related to the participation of the *Chalmers truck team* in the 2016 Grand cooperative driving challenge (GCDC). The GCDC was a part of i-GAME, a European research project that promoted the development of cooperative automated systems for vehicles. The event itself was a demonstration of cooperation between automated vehicles, but was also a competition between the participants to see which team could handle the different scenarios of the demonstration best [19]. The teams were judged not only on the individual achievements but also on cooperation and helpfulness, all to promote the progress of the field. The teams' HMI designs were also judged separately to promote the communication to other agents in the traffic [11].

The aim of this thesis was to use sensor fusion to enhance the performance of the automated system by increasing the precision and reliability of the data and by doing so give the team benefits in the GCDC. Sensor fusion is an umbrella term for different techniques whose common goal is to combine different sensor inputs to get an output of higher quality than from the individual inputs. One of the main tasks in this work was to enhance the accuracy of the positioning of the truck using data from GPS, Inertial Measurement Units (IMU:s), and actuators (steering wheel and gas pedal). This is a fairly common application of sensor fusion and a typical solution is to use a Kalman filter to combine sensor observations with calculated predictions. Using several different sensors to increase robustness has been proven effective [13], so furthermore this work will demonstrate how sensor fusion can be used to combine the information from laser range finders, sonars, and cameras to enhance the confidence levels of the localization of other cars and obstacles.

Additionally, as a secondary objective, this work also aimed to use the GCDC as a source of information and to gather data about which sensors and methods other teams were using. By doing this and by discussing the merits of different components and techniques the goal is to compile a summary on quality and usefulness of the contemporary resources used in the field of sensor fusion.

The level of success of the thesis work will be measured by comparing the results from test runs without any sensor fusion to test runs with different sets of sensor fusion techniques.

1.1 The scenarios

This section will briefly describe the different scenarios that constitutes the GCDC. Since the i-GAME project is all about cooperation between vehicles, these scenarios focus on situations where communicating vehicles might be better suited than a traditional autonomous vehicle that only react to passive information.

1.1.1 The lane merging scenario

This scenario starts with two lanes of traffic in the same direction, as can be seen in Fig. 1.1. At a certain point a message will be broadcasted that there is road maintenance up ahead on the road and that one of the lanes are blocked so that the two *platoons* will have to merge. The vehicles will then slow down and the ones in the blocked lane will each pair up with a vehicle behind them in the free lane. The vehicle in the free lane that is paired with the lead vehicle in the blocked lane will then create merging space in front of it by slowing down further. When the merging space is sufficiently large it sends a *safe to merge* message to the lead vehicle in the blocked lane after which it can safely join the remaining lane in front of the paired vehicle. The procedure can then be repeated with the new lead vehicle.

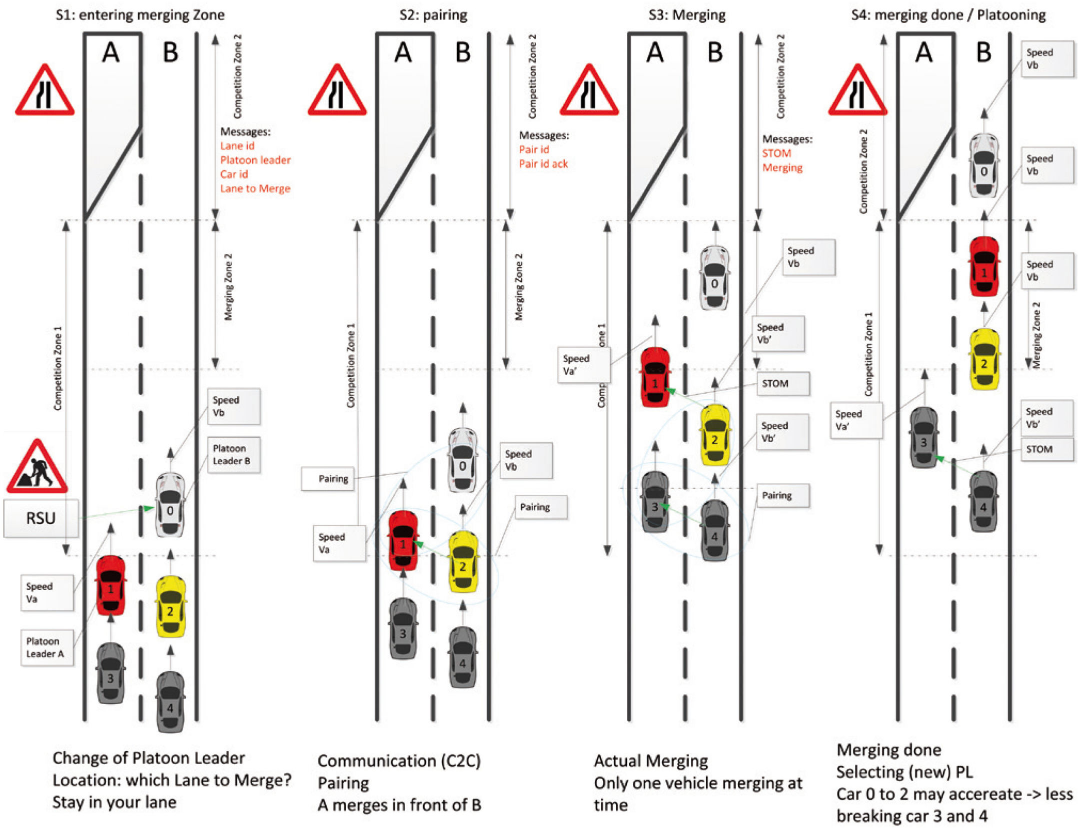

Figure 1.1: The lane merging scenario

1.1.2 The intersection scenario

Fig. 1.2 shows the intersection scenario, where two competing vehicles approaches an intersection from opposite directions while a reference vehicle approaches on an adjoining road. The competing vehicles should achieve a speed of at least 30 km/h before the area of the intersection, but as they get the signal that the reference vehicle approaches, they are supposed to adjust their speeds sufficiently for the other vehicle to join the traffic on the main road without interruptions.

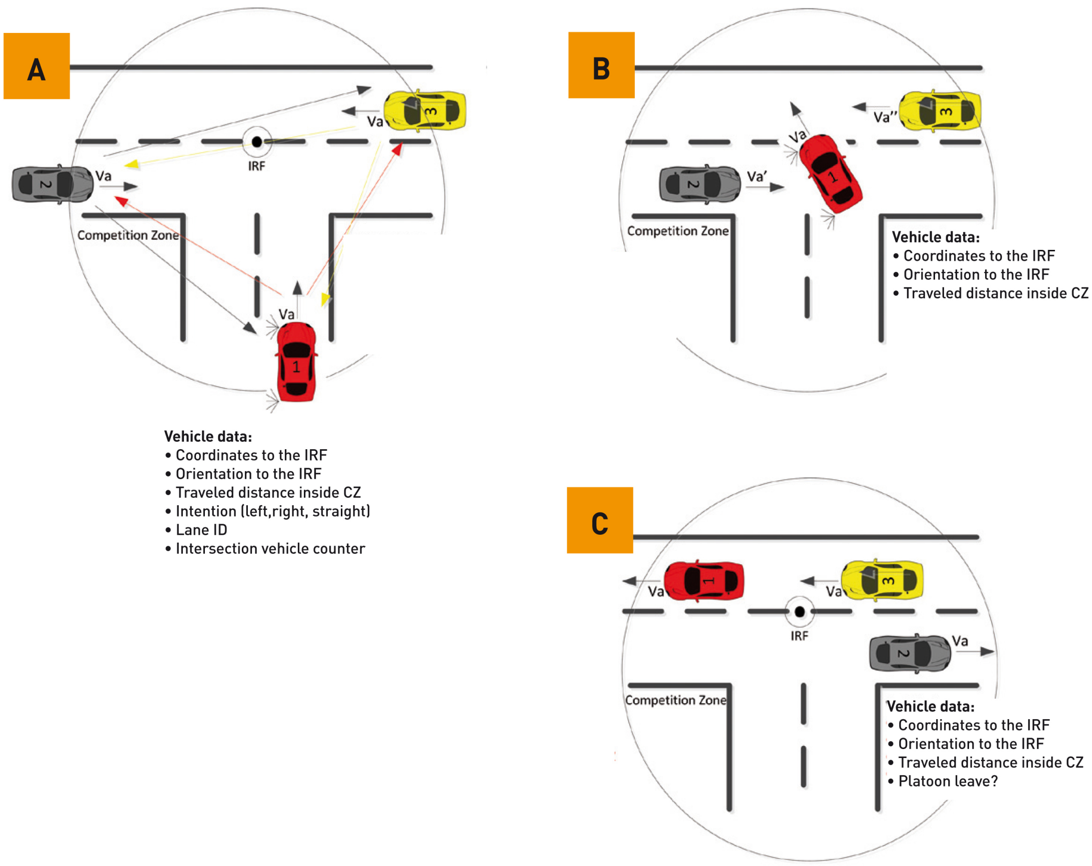

Figure 1.2: The intersection scenario

1.1.3 The emergency vehicle scenario

This was only a demonstration scenario to show how useful cooperation between autonomous vehicles can be in certain traffic situations and was not considered for competitive purposes. As can be seen in Fig. 1.3 the intelligent vehicles travels side by side on a two-lane road an emergency vehicle approaches from behind. As the autonomous vehicles receives an emergency signal from the Emergency Vehicle (EV) they are supposed to slow down and move to the sides of the road to make room for the emergency vehicle in the middle. When the EV has passed, the vehicles should resume their earlier speed and position.

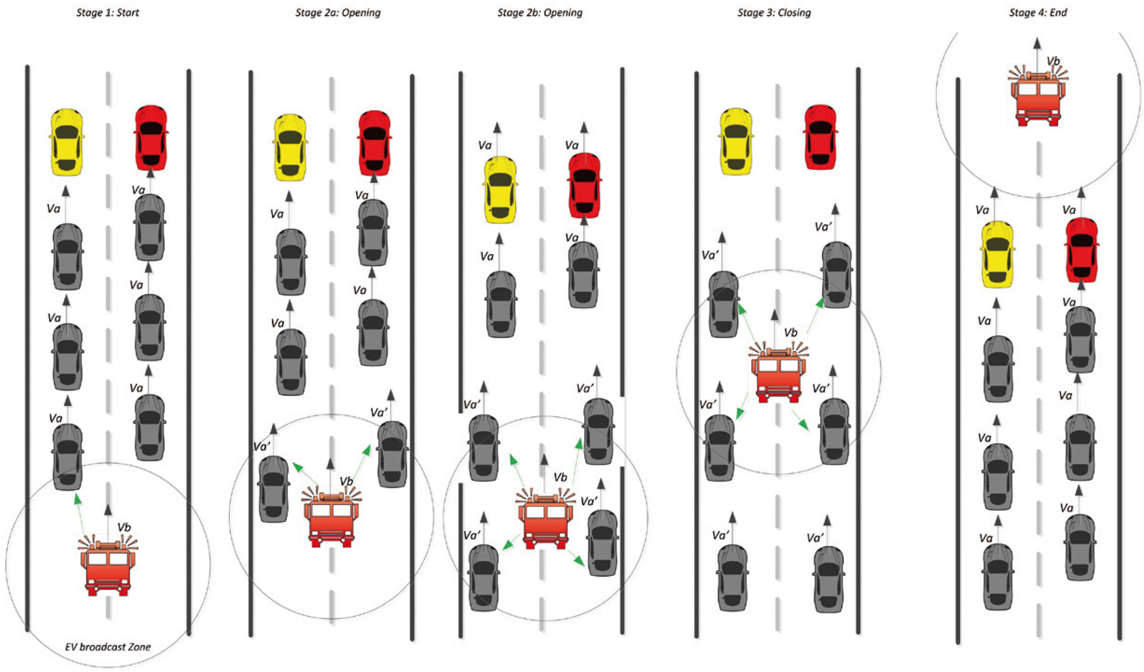

Figure 1.3: The emergency vehicle scenario

2

Background

2.1 Sensor fusion

A literal description of sensor fusion is that it is *the combining of sensory data or data derived from sensory data such that the resulting information is in some sense better than would be possible when these sources were used individually*. The concept has been around at least since the 1950s, even if the term ‘sensor fusion’ is a later construct [1]. In 1960 Kalman presented a paper explaining a mathematical algorithm that did just that [8]. What became known as the Kalman filter is still one of the most widely used methods in the field of sensor fusion today [1], albeit in some different forms than was originally presented [10][14], many of them developed by Kalman himself [9].

A Kalman filter basically consist of the two phases: Update and prediction. In the update phase all measurable state variables are updated with actual values from sensors. In the prediction phase a state model is used to predict a set of the state variables for the upcoming time step. In the next update phase the values from the sensors are merged with the predicted values to form the updated state variables. In this way the filter will balance the practical measurements from the sensors that might be unreliable against the theoretical values from the model that will never fit reality perfectly, to form a more stable interpretation of the situation at hand [1]. The area of motion tracking is especially adequate for the implementation of Kalman filters and has been used for that for as long as the filters have existed [17]. For a Kalman filter to work properly though, a very important choice is the choice of state model. For basic motion tracking a simple kinetic model might work, but a common choice for vehicle models is the planar bicycle model [15], as this model adds the benefit of taking the lateral forces on the tires of a vehicle into account [16].

While sensor fusion has been used in engineering since the 1950s [8], similar concepts have been used by evolution since time immemorial. Humans and most animals use the senses of sight, smell, hearing, touch and taste to perceive their surroundings, much like the sensors of an advanced piece of technology. Natures equivalent of a computer, the brain, then fuses the data from these sensors and produces a dynamic model of the world with the help of the previously obtained knowledge of the surroundings. One beautiful example of sensor fusion in the brain is the emergent property of distance judgment by combining the input from two (or more) eyes. In the end, the brain is able to predict the development of the environment and take informed decisions based on this knowledge.

So why is sensor fusion used? The reasons can basically be broken down into two groups: Limitations of the sensors and advantages of the algorithms. All sensors are to some degree subject to limitations in the following areas: Domain, frequency, propensity for breakdowns, precision, and certainty [1]. Many of these limitations can be remedied by adding more sensors without any advanced fusion of data, but especially imprecision and uncertainty can be greatly reduced by applying e.g. a Kalman filter. A good implementation of sensor fusion can fuse the data from several sensors to handle filtering of noise, produce emergent properties, and in some sense predict the near future.

Any system using sensors to measure the world can potentially benefit from sensor fusion and it is present in as diverse application as agriculture, home care and navigation [4][12]. As a result there seems to be a common belief that everything can be made better by the application of data refinement or by adding more data to the process. Some researchers are critical to this way of thinking though, and several papers are questioning the statement that sensor fusion always produces data of higher quality in some aspect [6], that the addition of data is always beneficial [5], and if the cost is worth the benefits [7]. This has of course in itself provoked the opposite side to try to prove that the algorithms provide inherent enhancements of the data [3]. The different examinations are of too limited scope to make a definite conclusion, but that there is cause for some healthy skepticism is evident.

2.2 Autonomous vehicles

In the vehicular industry, sensor fusion is making an impact on the area of active safety. The rapid approach of the milestone of self-driving vehicles makes the importance of active safety even greater, and the software behind active safety craves both accurate and reliable sensor data. The increase of computing power in vehicles and the relative lack of space for sensors makes the vehicle platform even more ideal for the application of sensor fusion.

The self-driving vehicle is something of the holy grail for modern technology and car and truck companies vie for the honor of being the first ones to present a fully autonomous and safe system. The level of automation in vehicles is usually measured using SAE's five levels of driving automation for on-road vehicles, see Fig. 2.1, from the J3016 standard. Even if some of the criteria of the these levels are somewhat subjective, the road to a level five vehicle seems long, but the truth is there are vehicles with level four automation driving around in real world traffic situations already today [20]. So far it is only in restricted areas and under certain conditions, but it still proves how far the field has come.

The idea of autonomous cars is not new; it has been around at least since the 1930s in science fiction novels and has appeared in many TV shows over the years. The first time that a fully autonomous car seemed not only a possibility, but a reality

SAE level	Name	Narrative Definition	Execution of Steering and Acceleration/ Deceleration	*Monitoring* of Driving Environment	Fallback Performance of *Dynamic Driving Task*	System Capability *(Driving Modes)*
***Human driver* monitors the driving environment**						
0	No Automation	the full-time performance by the *human driver* of all aspects of the *dynamic driving task*, even when enhanced by warning or intervention systems	Human driver	Human driver	Human driver	n/a
1	Driver Assistance	the *driving mode*-specific execution by a driver assistance system of either steering or acceleration/deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	Human driver and system	Human driver	Human driver	Some driving modes
2	Partial Automation	the *driving mode*-specific execution by one or more driver assistance systems of both steering and acceleration/ deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	**System**	Human driver	Human driver	Some driving modes
***Automated driving system* ("system") monitors the driving environment**						
3	Conditional Automation	the *driving mode*-specific performance by an *automated driving system* of all aspects of the dynamic driving task with the expectation that the *human driver* will respond appropriately to a *request to intervene*	System	**System**	Human driver	Some driving modes
4	High Automation	the *driving mode*-specific performance by an automated driving system of all aspects of the *dynamic driving task*, even if a *human driver* does not respond appropriately to a *request to intervene*	System	System	**System**	Some driving modes
5	Full Automation	the full-time performance by an *automated driving system* of all aspects of the *dynamic driving task* under all roadway and environmental conditions that can be managed by a *human driver*	System	System	System	**All driving modes**

Figure 2.1: SAE International's standard for vehicle automation levels

though, was at the third DARPA grand challenge (also known as the DARPA urban challenge) in November 2007 [21]. In this challenge several teams managed to get their vehicles to maneuver through traffic consisting of other vehicles and perform advanced tasks, such as merging and negotiating intersections. The DARPA urban challenge of course inspired other projects, such as the 2011 Grand cooperative driving challenge in Eindhoven, the Netherlands. This challenge was aimed to drive the development of intervehicle communication and cooperative autonomous driving and was designed as a competition between different research groups [18]. The outcomes of the first GCDC was both promising — vehicles of several different brands had cooperated successfully due to common communication standards, which had never been done before — and lacking — the trustworthiness of the data and the adherence to said standards was atrocious in some cases [18]. Therefore, a second GCDC event was eventually run within the 2016 i-GAME project [19], a framework in which this thesis was created.

3 Method

3.1 OpenDLV

The work on sension fusion performed in this masters thesis is implemented in a framework OpenDLV. OpenDLV is an open source software project for self-driving vehicles that was initiated and is currently developed at Chalmers' vehicle laboratory Revere. The software is mainly written in C++ and is heavily distributed and modularized by the design principle of microservices, and developed with ease of deployment and distributed development in mind. On the behavioral level, it is biologically inspired and mimics the same principles as biological agents do when mapping sensor input to action.

3.2 Sensors

For the purpose of sensing the world around it the truck was equipped with several sensors:

- A 2D lidar mounted at the front of the truck below the windshield to measure accurate distances to objects in front of the truck
- Two cameras that were mounted in the lower corners of the windshield inside the cabin to localize the lanes and identify other vehicles
- A GPS that was placed in the cabin, with the antenna mounted on the roof of the cabin, to track the absolute position of the truck
- An IMU installed in the cabin to measure the accelerations in different directions
- Six short range sonars that were distributed around the sides of the truck to sense if other vehicles were present where the cameras and lidar couldn't detect them
- A vehicle-to-vehicle communication unit (V2V) for intervehicle communication

3.3 Position tracking

The position of the truck is measured by the GPS, but such a sensor might sometimes deviate and give unstable measurements. Therefore, to track the position accurately the GPS position could be combined with a predicted position derived from other sensor data such as velocity and actuation feedback from the steering wheel and gas pedal. Combining measurements with predicted data, with the help of a well designed state model, is exactly, as described above, what a Kalman filter is designed to do. Therefore, such a scheme, as illustrated in Fig. 3.1, was used to monitor and maintain the truck's position [14].

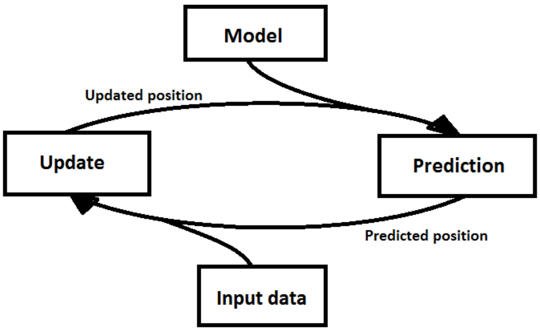

Figure 3.1: Block view of a Kalman filter

3.3.1 State models

At first a simple kinetic model with x and y-position, heading-angle θ as system state and longitudinal velocity v and the angle of the steering wheel ϕ as control input was used to model the system:

$$\text{State vector: } \begin{bmatrix} x \\ \dot{x} \\ y \\ \dot{y} \\ \theta \\ \dot{\theta} \end{bmatrix} \quad \text{Control vector: } \begin{bmatrix} v \\ \phi \end{bmatrix} \tag{3.1}$$

State transition function f:

$$\begin{aligned} \hat{x} &= x + \dot{x}\Delta t \\ \hat{\dot{x}} &= v\cos\theta \\ \hat{y} &= y + \dot{y}\Delta t \\ \hat{\dot{y}} &= v\sin\theta \\ \hat{\theta} &= \theta + \dot{\theta}\Delta t \\ \hat{\dot{\theta}} &= \tfrac{v\tan\phi}{l} \end{aligned} \tag{3.2}$$

The variable l is the distance between the front and rear wheel pairs.

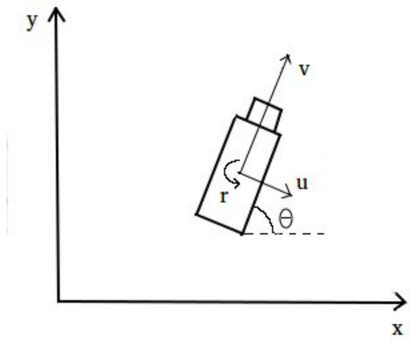

Figure 3.2: Outline of the truck variables

However, it was soon concluded that a more refined dynamic model was needed to accurately model the lateral movement. Therefore, a planar bicycle model (PBM) was instead used as a prediction model. The PBM models the lateral forces on the tires and it uses the lateral velocity and the yaw rate to calculate the lateral and angular accelerations, which can be seen in the following functions:

$$\begin{bmatrix} \dot{u} \\ \dot{r} \end{bmatrix} = \begin{bmatrix} \frac{-C_f - C_r}{mv} & -v + \frac{C_r b - C_f a}{mv} \\ \frac{C_r b - C_f a}{Iv} & \frac{-C_f a^2 - C_r b^2}{Iv} \end{bmatrix} \begin{bmatrix} u \\ r \end{bmatrix} + \begin{bmatrix} \frac{C_f}{m} \\ \frac{C_f a}{I} \end{bmatrix} \phi \tag{3.3}$$

Description of constants:

$$\begin{aligned} m =& \quad \text{mass of truck}(kg) \\ C_f =& \quad \text{front wheel cornering stiffness}(N/rad) \\ C_r =& \quad \text{rear wheel cornering stiffness}(N/rad) \\ a =& \quad \text{distance between front wheel and center of mass}(m) \\ b =& \quad \text{distance between rear wheel and center of mass}(m) \\ I =& \quad \text{moment of inertia}(m^2 kg) \end{aligned} \tag{3.4}$$

The lateral velocity is denoted u and the yaw rate is here denoted r, as shown in Fig. 3.2. These variables were added to the control vector, and since it is preferred to keep everything in a global reference frame, the lateral acceleration is split up and added to the transition functions of the velocities in the x and y directions:

$$\text{State vector: } \begin{bmatrix} x \\ \dot{x} \\ y \\ \dot{y} \\ \theta \\ \dot{\theta} \end{bmatrix} \text{ Control vector: } \begin{bmatrix} v \\ u \\ \phi \\ r \end{bmatrix} \tag{3.5}$$

The new state transition function:

$$
\begin{aligned}
\hat{x} &= x + \dot{x}\Delta t \\
\hat{\dot{x}} &= v\cos\theta - \dot{u}\sin\theta\Delta t \\
\hat{y} &= y + \dot{y}\Delta t \\
\hat{\dot{y}} &= v\sin\theta + \dot{u}\cos\theta\Delta t \\
\hat{\theta} &= \theta + \dot{\theta}\Delta t \\
\hat{\dot{\theta}} &= \tfrac{v\tan\phi}{l} + \dot{r}\Delta t
\end{aligned}
\tag{3.6}
$$

The $\dot{u}$ and $\dot{r}$ is calculated from the PBM equations in every time step. It might be worthwhile to mention that both r and $\dot{\theta}$ refer to the yaw rate, but r refers to the measured yaw rate, whereas $\dot{\theta}$ refers to the predicted yaw rate.

Since the vehicle model is non-linear, an extended Kalman filter needs to be used in which the function derivatives are estimated in each time step by using the function's Jacobian matrix, given as:

$$
J_f = \begin{bmatrix}
1 & \Delta t & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -v\sin\theta - \dot{u}\cos\theta\Delta t & 0 \\
0 & 0 & 1 & \Delta t & 0 & 0 \\
0 & 0 & 0 & 0 & v\cos\theta - \dot{u}\sin\theta\Delta t & 0 \\
0 & 0 & 0 & 0 & 1 & \Delta t \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix}
\tag{3.7}
$$

3.4 Pitch

As the GPS antenna was mounted at the roof of the truck cabin, and since the cabin is suspended above the rest of the truck and thus will exhibit relative motion, the pitch of the truck must be taken into consideration when calculating the position of the truck. By using a simple model of the cabin suspended upon two equal sets of spring and damper (see Fig. 3.3), the following equation can be derived:

$$
m\ddot{z} = -\frac{am\dot{v}}{b} - kz - d\dot{z} \tag{3.8}
$$

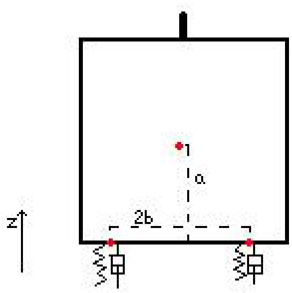

Figure 3.3: Model of the cabin suspension

Here a is the distance from the bottom of the cabin to its center of mass, b half the distance between the springs, k a spring constant, d a dampening constant and z the altitudinal displacement of the cabin above one of the spring and damper pairs. The longitudinal displacement l of the GPS antennae can be computed as follows:

$$l = -\frac{2az}{b} = \frac{a}{bk}\left(m\ddot{z} + \frac{am\dot{v}}{b} + d\dot{z}\right) \tag{3.9}$$

From here, there are three options of how to calculate the antenna displacement:

The first option is to obtain all the variables in Eq. 3.9. The IMU in the cabin directly gives the accelerations in the longitudinal and vertical directions ($\dot{v}$ and $\ddot{z}$) and the vertical velocity $\dot{z}$ can be estimated from the pitch rate which is also obtainable from the IMU. This method however uses several measurements and error-prone estimations in an already heavily approximated model.

The second option is to simply take the difference of the pitch measurements from the IMUs of the cabin and truck base respectively and calculate the displacement trigonometrically. This might not be the most accurate way, but it is simple and straightforward although it requires some tuning.

A third option is to make the assumption that since the displacement is primarily caused by the longitudinal acceleration it is possible to make an accurate approximation using only the acceleration as input. One way of implementing that is illustrated in Eq. 3.10.

$$\begin{aligned} \dot{x} = &\quad (v + p\dot{v})\cos\theta - \dot{u}\sin\theta\Delta t \\ \dot{y} = &\quad (v + p\dot{v})\sin\theta + \dot{u}\cos\theta\Delta t \end{aligned} \tag{3.10}$$

Here p is the unknown constant that will approximate the complex system, and thus has to be estimated. A good estimation might be the inverse of the steady-state frequency of Eq. 3.9, i.e. $\sqrt{m/k}$.

3.5 World awareness

For any intelligent agent to be able to take an enlightened action that interacts with its environment, such as moving in the case of a vehicle, it has to know its surroundings. To keep updated about the world around it, the agent needs sensors to interpret and maintain an internal description of the external environment [10]. In OpenDLV the world awareness is split into two parts: One that takes care of the vehicles position relative to the world, and one that takes care of the vehicles immediate surrounding.

The function maintaining the trucks own position is called `geolocation` and maps the trucks location to absolute longitudinal and latitudinal coordinates. This is done by applying the Kalman filter and state models from Sect. 3.3.1 on the data supplied from the GPSs and IMUs.

The immediate surroundings is taken care of in the function `scene`. The state of the external environment is described by a set of objects and surfaces that are maintained and updated by the sensors. An object is in this case most often a car, but might also be an obstacle such as a wall or a tree. The object is primarily characterized by a direction, a distance and an angular size, but can also contain information about object type, angular rate (a measure of relative movement), angular size rate and information sources. A surface in OpenDLV usually depicts a lane or road and is defined by a quadrangle described by its four corners and if it is traversable or not. Additionally it can contain information about information sources, other adjoining surfaces and whether it is possible to move to these.

Just as an organism trusts its different senses to varying extents (like how a human in most instances trusts the eyes over the ears), OpenDLV is supposed to have nuanced faith in its sensors. For this reason, a confidence system for the different information sources and data types is implemented in `scene`. Here each attribute is assigned a confidence value between 0 and 1 for various degrees of certainty or -1 if there is no information. These values are then updated as new information is added by the sensors and as time passes by. To enhance the probability of an enlightened decision all objects and surfaces are thus combined with a confidence level and additionally contains a unique confidence value for every listed property.

3.6 Interviews

In an effort to evaluate which sensors and sensor fusion techniques that were used in GCDC a list of questions were compiled to serve as a base for a series of interviews of the competing teams that was conducted during the competition:

- Which kinds of sensors is used on the vehicle?
- Was there any plans on using more or other sensors on the vehicle?
- How is the sensor data processed?
- What data is used to determine the position of other vehicles?
- What would have been the next improvement if there were more time?

4 Results

4.1 Position tracking

To evaluate the Kalman filter with the kinematic model a set of test data with noisy GPS data was run through the filter. The result of this can be seen in Fig. 4.1

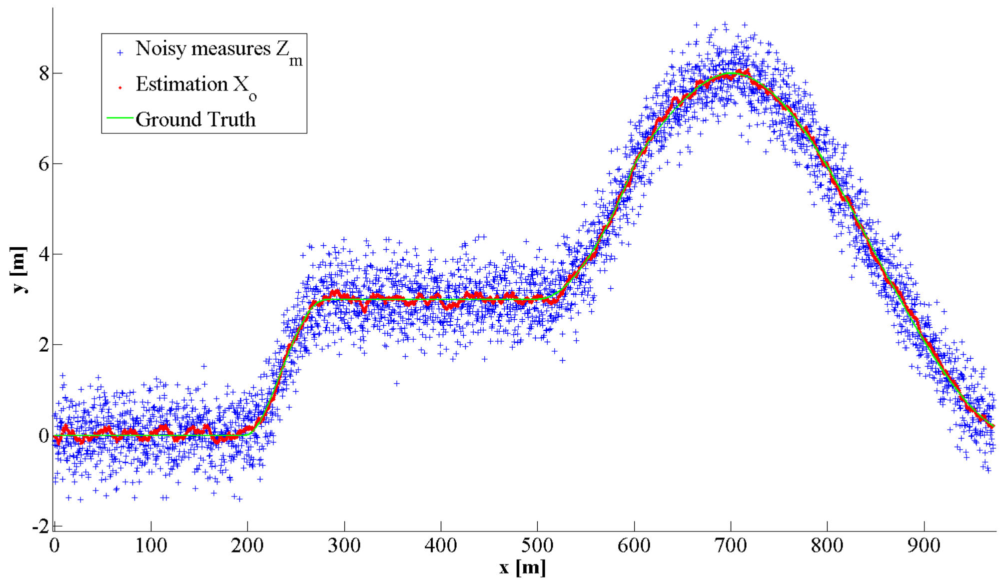

Figure 4.1: Evaluation of Kalman filter using noisy test data

After establishing that the kinematic model was suitable it was also evaluated on real data from a test drive at Lindholmen, Göteborg. The results from this can be seen in Fig. 4.2

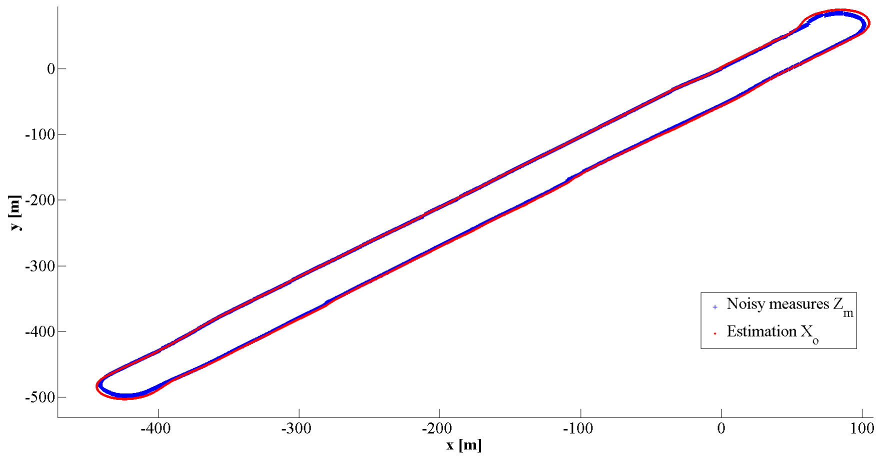

Figure 4.2: Evaluation of the kinematic model using real data

As can be seen in Fig. 4.2 the kinematic model leaves room for improvement and so it was complemented with a dynamic bicycle model.

To see how robust the model would be and how it responds to data loss it was also tested on data with missing GPS points. An evaluation of the response to data loss in the dynamic model can be seen in Fig. 4.3

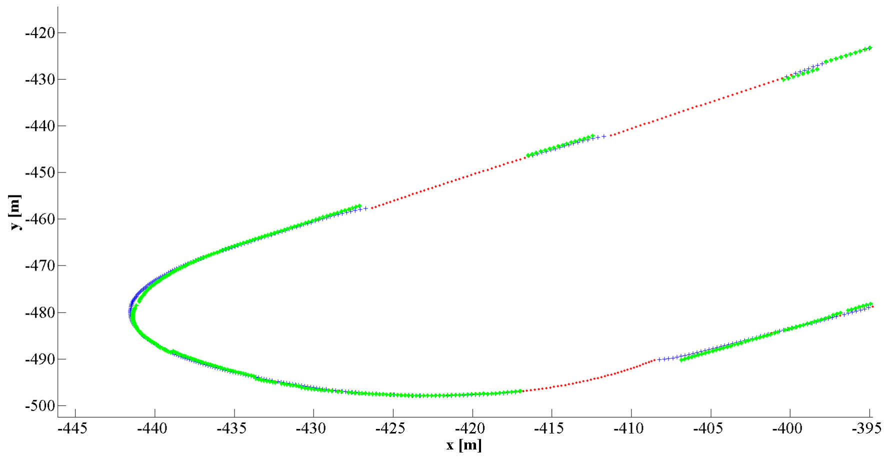

Figure 4.3: Evaluation of the dynamic model with induced data loss. The green blobs represent real GPS data, the blue pluses represent estimated position from the GPS data and the red dots represents the predicted position when no GPS data is available

4.2 Pitch compensation

To evaluate the pitch compensation schemes two GPSs were needed. The one on the roof of the cabin used under normal circumstances and one mounted on the frame of the truck, to compare with to estimate the relative dislocation. The truck was driven with this setup for several runs and a rendition of one of these runs can be seen in Fig. 4.4.

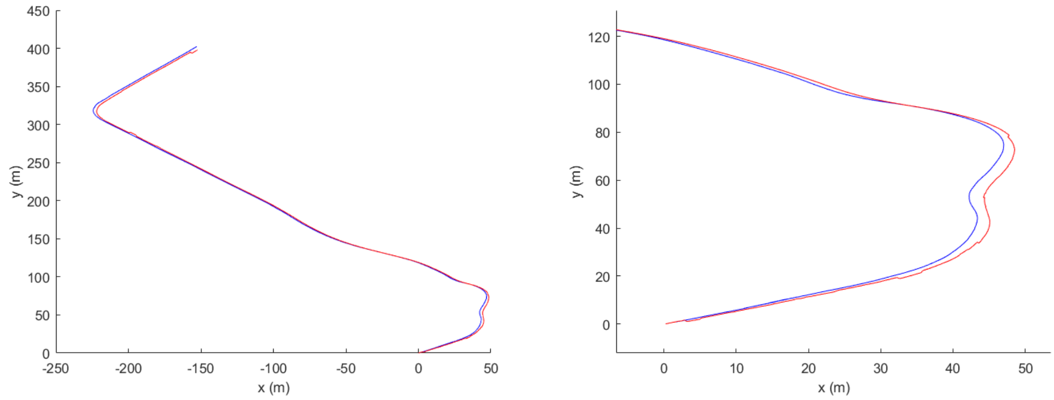

Figure 4.4: Visualization of a test run with two GPSs mounted on the truck which shows the path driven by the truck. The figure to the right is a zoom-in of the first part of the drive.

As can be seen the two GPSs deviate quite far from each other, and notably in a lateral direction which (hopefully) does not reflect reality as that would mean that truck was was travelling side-wise. When considering the absolute distance between the sensors in the same run — which is depicted in Fig. 4.5 — it can be observed that it varies by more than 4 meters (and in some runs up to 10 meters).

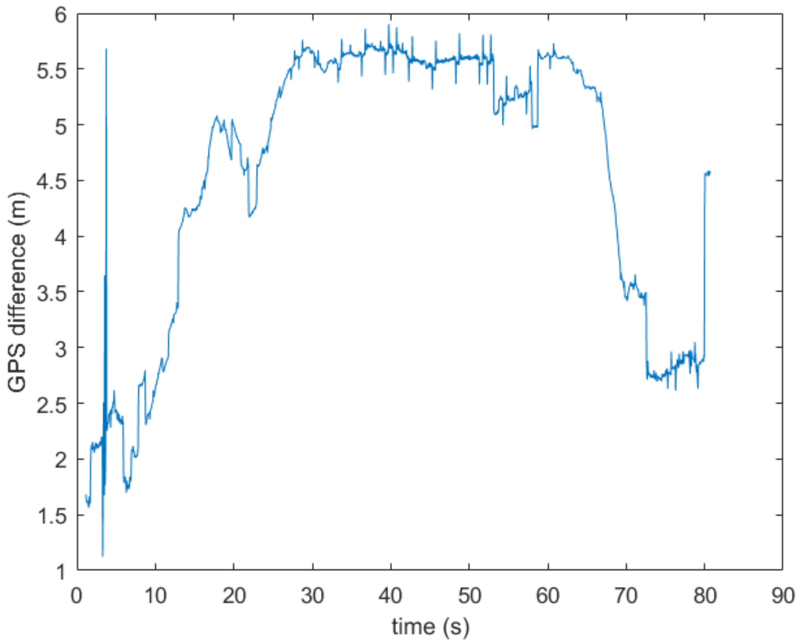

Figure 4.5: Absolute distance between the two GPSs

As the expected variation is less than a meter this data can not be used to evaluate any pitch compensation scheme. A rough estimation of the second compensation scheme from Sect. 3.4 done on data from the same run shows feasible values in Fig. 4.6 which confirms that it is not due to simple conversion errors.

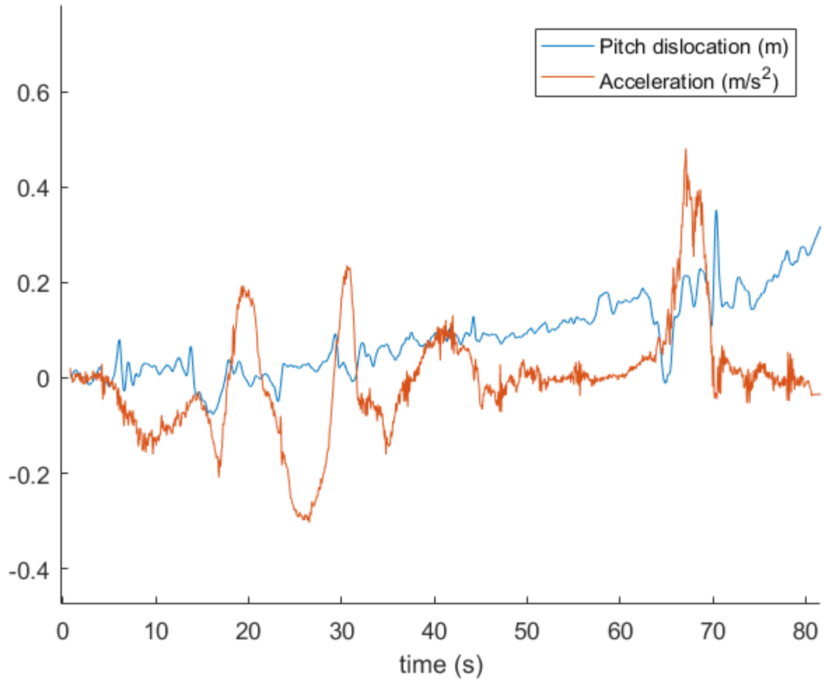

Figure 4.6: Rough estimation of a pitch compensation scheme for the dislocation of the main GPS receiver versus the acceleration of the truck

Despite some drifting — that is probably caused by the rough estimation — it can be seen that the more significant deviations in dislocation corresponds to similar deviations in the acceleration. Since the deviations of the GPSs was of magnitudes larger than the pitch compensation no further evaluation was relevant, but with more accurate sensor data a pitch compensation scheme would have been a valuable tool to increase precision.

4.3 World awareness

In the end the function `scene` was used in a very limited scope, only to measure the distance to the closest vehicle in front of the truck and keep track of the lanes. A lot of work and thought went into implementing the confidence system though and to try to visualize the surroundings of the truck in a good way. An early example of a visualization scheme based on data from the lidar can be seen in Fig. 4.7

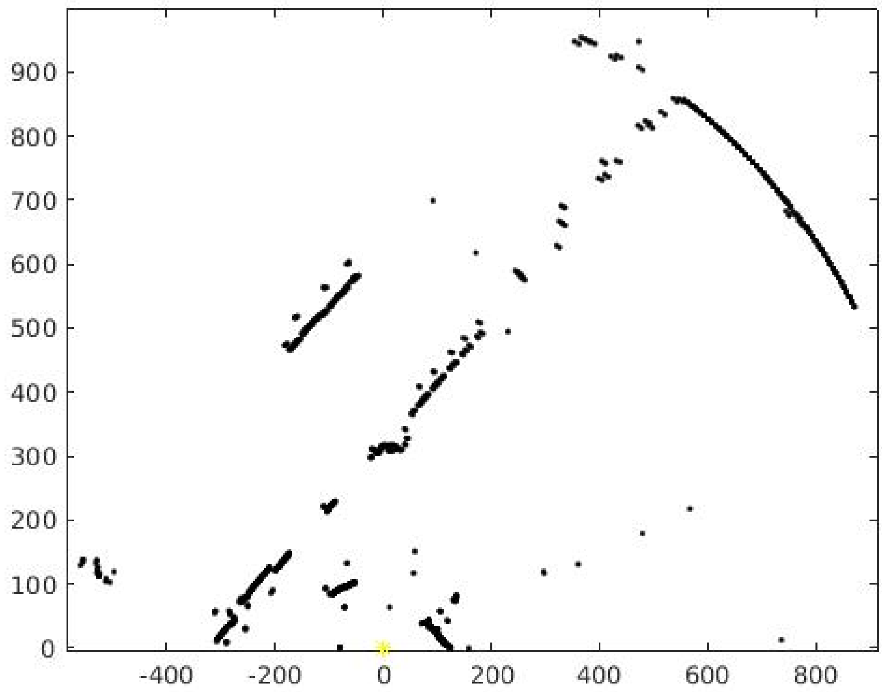

Figure 4.7: Visualization of a lidar scan

4.3.1 Confidence

Although the confidence scheme was implemented and used to some extent, time did not permit any tuning or testing of different thresholds and parameters, so the system was mostly running on default values. Even so the process of setting it up resulted in a lot of insights.

The information on an object can originate from several different sources in the OpenDLV system and can therefore contain very diverse information of various degrees of certainty. The lidar is one of the most trusted sensors as it gives very accurate information, but does not provide information for all properties. It provides very reliable data on distance, direction and angular size of an object, but can tell nothing about the type of object it detects.

Information from the cameras via vehicle detection on the other hand gives pretty good information on whether the object is a vehicle or not, while the information on direction and angular size is a bit less trusted and the distance even less so. The cameras also gives information on surfaces via lane detection and since it is the only source of surfaces so far its relative confidence is rather high.

The V2V unit gives perfect information on the object type, but depending on who the source of the signal is, the quality of the rest of the information provided might range from excellent to worthless. For the scope of the GCDC there were therefore plans of evaluating the different teams and assigning different levels of confidence according to their performance, but the probability of managing something similar in a real life application is low, at least in a foreseeable future.

The sonars are in a way the least trusted of the sensors as they give just one measuring point, that on top of that is pretty fuzzy. The purpose of them though is only to tell if there are any objects at the vehicles sides and not to give an exact position, and with the deployment of several at each side they still fill their function.

As the surroundings of a vehicle on the move are highly dynamical all new information has to be matched and combined with the environment state from the previous time step to maintain a reliable environment description. For an object, the matching is done by comparing its direction and distance to these of the objects saved in the environment function `scene`, with a slight emphasis on the angle, since the forward vehicles are of most interest. If there is no match a new object is created with the information from the sensor and relatively low confidence depending on the source sensor. If there is a match however, the properties of the existing object in `scene` is updated according to the quality of the match and the confidence of the properties of the two objects. The updating process might also produce new properties such as angular rate and angular size rate, which are not actually measured by the sensors. Each property has its own way of updating confidence, but some general things might be said about the update process:

- A new matching reading increases the confidence of the object and generally increase the confidence of properties
- Most property confidences increase with different amounts according to sensor types and how well the properties of the objects matches up
- Detection of the same object by a new sensor gives a confidence boost
- Distance measures are updated with a bias for closer readings for safety reasons
- Property values are updated as a weighted average between the previous and new values weighted with a bias for high confidence and a heavy bias for new data, due to the dynamic nature of the system

Surfaces are mainly matched by overlap percentage, but otherwise the process is much the same as for objects.

Since no sensors are fully reliable, objects and surfaces have a fade time —- a period of time in which it doesn't disappear even if no new detections of it are made by the sensors. This is a safety feature, so that the vehicle does not ignore another vehicle or obstacle just because the sensors missed it for a few moments. As time passes by the reliability of the information deteriorates because of the dynamic nature of the system, and so the confidence levels of the object and surface properties decrease with time.

4.4 Interviews

In an effort to evaluate which sensors and sensor fusion techniques was used in GCDC a series of interviews was conducted with members of the attending teams. In most cases only one or two of the team members were interviewed and the ones interviewed was not always the expert in the team on the topics at hand. Even though effort was taken to find the most knowledgeable member in the areas of sensors and sensor fusion in each team they were not always available due to the limited time at the site. This might of course skew the results somewhat, but only responses where the interviewee was reasonably sure have been recorded. An extensive compilation of the results of the interviews can be seen in appendix A.1, but short summaries of the results will be provided below.

4.4.1 Sensors

All the participating teams use one or more GPSs to localize the vehicle, usually mounted on the top of the vehicle and of course all teams also used the mandatory V2V to communicate with the other participating vehicles. All competitors used data collected from actuation feedback, but maybe somewhat surprising more than one team had to use their own solution to measure the actuation, since they did not have access to the signals from the inbuilt sensors. All teams had planned to use either a lidar or a radar as well, but a few of the teams either did not manage to incorporate their sensor into the SW or decided to be content with the data from the V2V units to localize other vehicles. Almost all teams had also equipped their vehicle with cameras, often more than one, but very few of them actually used them in the competition setup.

4.4.2 Sensor fusion

Most teams used Kalman filters to determine their own position and some few also kept track of the other vehicles using Kalman filters. Most of the teams had thoughts on using confidence based merging of objects and properties although few used it and those implementations ranged from simple mean values to advanced models such as regressive object fusion. Some teams also used sensor fusion to keep track of the environments: The KTH car team was setup to use simultaneous localization and mapping (SLAM) to keep track of their surroundings and the German team had scanned the environment prior to the competition day to be able to use the pre-processed map in their environment scheme.

4.4.3 Identification and localization of other vehicles

All teams used the information provided by their V2V units, but the level of trust in the information given ranges from teams that trusts the information fully and only uses that for localization to teams that only use the V2V information for identification and trust other sensors to make the localization. A few teams even assigned different confidence to data from different teams according to how reliable that information was. Many of the teams kept track of the other teams positions in their environment models and some of the teams even did it by maintaining state vectors updated by their own Kalman filters.

4.4.4 Unused equipment and techniques

Virtually all of the competing teams had equipment and techniques that they did not use either because of technical issues or because of lack of time, ranging from simple sensors that were incompatible to classification schemes that proved to be too computationally heavy to be of any use.

5

Discussion

5.1 Evaluation of the methods

In this section the sensor fusion methods used for Chalmers truck team's participation in 2016 GCDC will be evaluated and discussed.

5.1.1 Localization

The work on position tracking and implementing the Kalman filter was of utmost importance for the project. A trusted and accurate location is a fundamental base for any autonomous agent in traffic and without it all other work would have been of little consequence. In the end the implemented models fell out well and met the requirements of the competition without fault. In future applications of the truck software the localization schemes will have to be improved though, to counter a tendency for drifting in curvy scenarios.

5.1.2 Pitch compensation

The pitch compensation schemes never played any part in the competition as the benefits were to small to be relevant at that time. The reason for this was due to the fact that the speeds (and thus accelerations) in the scenarios were relatively low and the requirements on accuracy relatively lax, but had the rules been as strict as initially proposed (with a required accuracy of under a decimeter) a compensation scheme would definitely have been needed.

After the competition the work on the pitch compensation scheme continued for a while and data was gathered to evaluate the performance of the different schemes. Unfortunately the data gathered was of to low precision to show anything more than a vague hint of that the schemes would be useful on data of higher quality. Since the evaluation data was gathered via an extra GPS mounted on the chassis frame a feasible real world solution could be to compensate with the support of such a sensor even though that second GPS might be more prone to data loss due to the location. In a commercial setting however, every penny counts, so if a compensation scheme could replace an extra sensor, that would be highly desired and therefore there might be a future for such a scheme anyway.

5.1.3 World awareness

In the version of the code used in the competition the world awareness was reduced to keeping track of the lanes and the vehicle in front of the truck. This proved to be enough from a competitive point of view, but it leaves a lot of potential for future improvement. For keeping track of one lane and one vehicle the developed confidence schemes proved a bit overdeveloped, but it worked well in the context of the competition and with adjusted parameters it should be able to handle larger scopes as well. That would of course require a lot of work regarding tuning and optimization of factors and thresholds, but with the emergence of the fields of big data and neural networks the task of such parameter optimization should prove more feasible in the near future.

Even if a lot of the work on world awareness never made it to the version of the software that partook in the competition, it still gave some interesting insights into the problems that the automotive industry faces as the implementation of fully autonomous vehicles draws near. The work on confidence was inspiring and gave an indication on how wide the range of areas still in need of improvement is and just dipping a toe in the seemingly unending sea that is object identification via neural network classification of visual data shows just how daunting the specific area of visual recognition is. The work on visualizing the intent of the truck to both driver/passenger and surroundings will also be interesting and especially possible to follow even for a layman in the years to come.

5.2 Evaluation of the GCDC

This section contains evaluation and discussion on the impacts the 2016 GCDC had and will have on development of the areas of sensors and sensor fusion.

5.2.1 Sensors

The fact that all teams used a V2V unit is no surprise since the GCDC was a project aimed at improving cooperative behaviour and thus required one. The fact that the different teams placed a very varied amount on trust in the data projected by the V2V units gives a somewhat bleak outlook on their use in real world applications, which is further discussed in Sect. 5.2.4.

The same can be said about GPSs: Since each vehicle had to broadcast their position continuously it is really no surprise that all teams used GPS. A position can be obtained through other means (knowing the start position and using odometers and IMUs etc as an example), but the GPS solution is by far the easiest and most effective solution. Most teams were satisfied with the accuracy of their GPSs and from the context of the competition there is nothing that indicates that the emergence of autonomous vehicles will trigger any undue development in this field.

The fact that all teams used actuation feedback is also a expected: It is much easier to control your own state if you know what is being done to change it. When it comes to the use of lidar or radar all teams seems to agree that one of them should be used for precise distance measuring and that they are mutually redundant. Which one to use seems to be up to personal taste and ease of access, although the teams with access to both lidar and radar preferred the radar due to its slightly higher confidence. In this area most teams felt that there were room for improvement and the development of self-driving vehicles will most likely spur an even more intense evolution of range finding sensors.

As it comes to cameras all teams were very positive to the use of them, although very few actually used them and those who did so only used them for verification. This might very well be because all the vehicles in the competition were equipped with V2V units and thus could find all relevant vehicles that way, and since the scenarios required no identification of other objects, such as lanes or obstacles. In a more realistic environment cameras would be necessary and heavily relied upon to provide a wider range of information. The quality of the cameras used were not a topic of any noteworthy interest as it seemed that even cameras of quite low quality was able to provide the data necessary for the tasks in question. There is undoubtedly a need for high quality camera equipment in the future of autonomous vehicles, but the main focus will be on the software handling the data they produce. The process of producing good data and training networks handling identification of cars, trucks and other objects of interest is a daunting project and even with the rapid development in the field of big data this will be an area of interest for many years to come.

A final observation to remark upon regarding the sensors in the competition is the fact that none of the other teams had sensors in any other direction than forwards. This can with high confidence also be blamed on the somewhat simplified nature of the scenarios. In the real world, sensors in all directions will be needed to create a holistic view of the vehicles surroundings.

5.2.2 Sensor fusion

The usage of sensor fusion by the teams in the GCDC ranges from none at all to sophisticated confidence systems. In particular most teams used some variant of the Kalman filter for localization which will be further discussed in Sect. 5.2.3. Beyond that sensor fusion was somewhat scarce in the competition and limited to different classification and confidence schemes. One reason for this might be the fact that most teams used relatively few sensors, which once again might be because of the layout of the requirements of the scenarios. Since the scenarios lacked many of the challenges one might encounter in real life, the vehicles did not need that many sensors and with fewer sensors the need for sensor fusion naturally decreases. With the aim that the vehicles should be able to make their own way in the real world, the need for more sensors in more directions will rise and so will the need for robust sensor fusion.

5.2.3 Kalman filters

The scientific community as a whole seems to a agree that one of the variations of Kalman filters is they way to go for localization using GPS data, possibly backed up by an IMU [1].The experiences and interviews from the GCDC only reinforces this, as all teams that used sensor fusion for positioning did so using Kalman filters. However, some of the teams didn't use any sensor fusion at all since the raw data received from their GPS were enough in their opinion. All equipment has flaws though, and the possibility that there will be situations where data loss and low information quality occurs is very likely and at least for these reasons a Kalman filter with a strong state model will always be a great complement to good sensors. How advanced the state model of a filter needs to be, is up to the architect — a simple kinetic model in a regular Kalman filter is enough to counter light data loss and normal noise, while an advanced dynamic model in an extended or even unscented Kalman filter may compensate for poor sensor quality or severe data loss.

As for the specific case of the Chalmers truck team's participation in the GCDC it ended up with the use of an extended Kalman filter on a simple kinematic model. This was mostly because of the fact that a more dynamic model (such as the lateral bicycle model introduced in Sect. 3.3.1) would have needed data from an IMU and since the team didn't get the time to integrate the IMUs into the system that became impossible. However, since the scenarios in the GCDC did not call for much movement in the lateral direction, this was not a problem for the competition, but for future use of the truck in more complex areas the it is highly recommended to use an implementation of the developed dynamic model.

5.2.4 Cooperative autonomous vehicular systems

The thought of cooperative autonomous units is very appealing on a visionary plane, but there are a lot of different companies operating in the vehicular industry and aspiring on presenting their own autonomous vehicles, and getting them to cooperate is a whole other story. There is simply no innate incentive in being the first to develop a cooperative system that your competitors can use for free, and charging for it would most likely cause them to develop their own, which would defeat the purpose with a unified cooperative system. The probability that they will agree on a standard with acceptable accuracy on their own is therefore very low without at at least some incentive from one authority or another. The most feasible solution would be for the authority to develop and govern the whole cooperation part of the system themselves, since that could also lead to some more abstract, secondary benefits in areas like traffic flow and city planning.

The GCDC was hosted by i-GAME which is an EU project, and EU is an example such an authority. It is however also a very slow and bureaucratic authority and the development of automated solutions in the vehicular industry is currently very rapid, so the question is if it will be able to reach any conclusions regarding cooperative autonomous vehicular systems in time for them to be of any use, but that is for the future to tell.

6

Conclusion

That the future will contain self-driving vehicles is certain, which puts high demands on the technology involved. That the system controlling the truck needs to be able to handle situations with malfunctioning sensors, uncertain data and corrupt information is evident if the autonomous vehicles shall ever be safe enough for society to accept them. This puts strain on the requirements for both sensors and algorithms and so the need for qualitative sensor fusion will prevail for an unforeseeable future.

The GCDC is a great incentive for the development of systems capable of handling complex traffic situations. For obvious reasons the scope of the scenarios needed to be limited though and the main focus of the event was to further the cooperation between autonomous vehicles, which actually unburdens the need of sensor fusion in some sense. Disregarding the reasons the 2016 GCDC did not give the participating teams enough incentive to equip the vehicles with all the sensors that a self-driving truck would need in a real traffic situation. As a result of that the usage of sensor fusion was limited overall, but some conclusions can be drawn anyway.

The usage of sensor fusion was obviously not a necessity for participating in the competition, since some of the teams participated with good results without any sensor fusion at all. However, the fact that the winners, Halmstad University, used both Kalman filters and confidence schemes gives hope for a development in the right direction. If it was an advantage with sensor fusion in a competition with limited scope, where some teams' success lived and died on the premise of accurate sensors, it is evident that it would be mandatory for any real world application. The implementation of that sensor fusion is less defined though. The development of new, or fine-tuning of old algorithms and determining of appropriate parameter values will be a huge undertaking in the years to come.

Bibliography

[1] W. Elmenreich. An Introduction to Sensor Fusion. *Research Report 47/2001, Vienna University of Technology*, Austria, 2002

[2] P. L. Rothman and R. V. Denton. Fusion or confusion: Knowledge or nonsense? *SPIE Data Structures and Target Classification*, 1470:2–12, 1991

[3] P. J. Nahin and J. L. Pokoski. NCTR plus sensor fusion equals IFFN or can two plus two equal five? *IEEE Transactions on Aerospace and Electronic Systems*, 16(3):320–337, 1980

[4] C. Thomas. Sensor Fusion, Foundation and Applications. ISBN: 978-953-307-446-7, 2011

[5] B. V. Dasarathy. More the merrier ... or is it? - sensor suite augmentation benefits assessment. *In Proceedings of the 3rd International Conference on Information Fusion*, volume 2, pages 20–25, France, 2000

[6] C. A. Fowler. Comments on the cost and performance of military systems.*IEEE Transactions on Aerospace and Electronic Systems*, 15:2–10, 1979.

[7] R. R. Tenney and N. R. Sandell jr. Detection with distributed sensors. *IEEE Transactions on Aerospace and Electronic Systems*, 17(4):501–510, 1981.

[8] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Fluids Engineering* 82.1, pp. 35–45, 1960

[9] R. E. Kalman, R. Bucy. New Results in Linear Filtering and Prediction Theory. *Transaction of the ASME, Series D, Journal of Basic Engineering*, Vol 83, 1961

[10] J. L. Crowley. Principles and Techniques for Sensor Data Fusion. *Signal Processing* 32(1-2):5-27, 1993

[11] O. Benderius, C. Berger, V. Malmsten Lundgren. The Best Rated Human–Machine Interface Design for Autonomous Vehicles in the 2016 Grand Cooperative Driving Challenge. *IEEE Transactions on Intelligent Transportation Systems* 19(4), 2018

[12] J. Llinas and D. L. Hall. An introduction to multi-sensor data fusion. *Proceedings of the 1998 IEEE International Symposium on Circuits and Systems* 6:537–540, 1998

[13] M. Bertozzi et al. Obstacle detection and classification fusing radar and vision. *IEEE Intelligent Vehicles Symposium* pp. 608–613, 2008

[14] G. Welch and G. Bishop. An Introduction to the Kalman Filter. *University of North Carolina at Chapel Hill*, 2006

[15] S. Sen, S. Chakraborty, A. Sutradhar. Estimation of vehicle yaw rate and lateral motion for dynamic stability control using unscented Kalman filtering (UKF) *Michael Faraday IET International Summit 2015*, Paper ID: 68, 2015

[16] A. Turnwald, S. Liu. Adaptive Trajectory Tracking for a Planar Two-Wheeled Vehicle with Positive Trail *2018 IEEE Conference on Control Technology and Applications (CCTA)*, 2018

[17] T. Nishimura. Error bounds of continuous Kalman filters and the application to orbit determination problems *IEEE Transactions on Automatic Control*, Volume: 12, Issue: 3, 1967

[18] J. Ploeg, S. Shladover, H. Nijmeijer, N. van de Wouw. Introduction to the Special Issue on the 2011 Grand Cooperative Driving Challenge *IEEE Transactions on Intelligent Transportation Systems*, Volume: 13, Issue: 3, September 2012

[19] J. Ploeg et al. Cooperative Automated Maneuvering at the 2016 Grand Cooperative Driving Challenge *IEEE Transactions on Intelligent Transportation Systems* Volume: 19, Issue: 4, April 2018

[20] A. St. Antoine. Driverless Cars: Where They Stand Now *Automobile Magazine https://www.automobilemag.com/news/autonomy-levels-cars-driverless*, 4th January 2019

[21] DARPA Urban Challenge *https://www.darpa.mil/about-us/timeline/darpa-urban-challenge*

A

Appendix 1

A.1 Interviews

	What sensors do you use?	How do you process yor sensor data?	How do you determine the positions of other vehicles?	Addititional information, such as unused sensors and techniques
KTH Stockholm (Truck)	GPS, IMU, Radar, Camera, Actuation feedback, V2V	Kalman filter on GPS and IMU data for own position. Point masses in a Kalman filter to keep track of others	V2V and radar, trust radar more than others	Planned for relative corrections of the distances to others, but got no time
KTH Stockholm (Car)	GPS, IMU, Lidar, Camera, Actuation feedback, V2V	SLAM, Stochastic filtering	V2V and radar, trust lidar more than others	Did not participate due to mishaps with equipment
Heudiasyc (France)	GPS, IMU, Lidar, Actuation feedback, V2V	Minimal sensor processing	V2V and lidar, but don't trust others	Had cameras, radar and more advanced lidars that were not used
AnnieWay (Germany)	GPS, IMU, Lidar, Camera, Actuation feedback, V2V	Regressive object fusion. Extended Kalman filter on bicycle model for positioning. Mapping coordinates to preprocessed map	Rely on th V2V, but validates with lidar and camera	Loads of extracted data that weren't used
University of Latvia	GPS, Actuation feedback, V2V	Use only raw data; No processing	Trust everyone via V2V	Had lidar and cameras that couldn't be hooked up to current system
Halmstad	GPS, Radar, Actuation feedback, V2V	Kalman filter on GPS for own position and added radar for others'. Confidence system to evaluate the precision of information	Puts relative trust in sensors and other teams via V2V based on a confidence based system	Used built in cruise control for speed
Drivertive (Spain)	GPS, Radar, Actuation feedback, V2V	Kalman filter for postion and object merging to keep track of others	Rely on sensors and uses V2V for validation and additional information	
Chalmers Truck Team	GPS, Actuation feedback, V2V	Kalman filter for position	Rely on V2V messages from others to determine their position	IMU, Cameras, sonars and lidar with respective systems for confidence, merging and validation weren't used
Benchmark vehicles of the organizers	GPS, IMU, Radar, Camera, Actuation feedback, V2V	Kalman filter for position and object merging to keep track of others	Trust others, but validates with camera	

Note: The list of used components and systems is in no way exact as it depends on the knowledge of the individual interviewed team members

Figure A.1: Interviews with members of the teams attending GCDC regarding sensors and sensor fusion