Advances in Multi-Sensor Data Fusion for Ubiquitous Positioning: Novel Approaches for Robust Localization and Mapping

Patrick Robertson, Matthias Röckl, Michael Angermann,
Institute for Communications and Navigation, German Aerospace Center (DLR),
Oberpfaffenhofen, Germany; firstname.lastname@dlr.de

Abstract

In this paper we argue why robust positioning in transportation applications is best achieved by multi-sensor fusion. Furthermore, we suggest that sensor fusion processing be performed in a probabilistic fashion and that in the majority of relevant practical applications one should draw on utility theory in order to make decisions that will be of the highest expected benefit given the current circumstances. Simply stated, it is a fact that all sensors are prone to errors or failure. Only if we model these errors correctly, and account for all possible failure modes, are we able to implement systems that reap the benefits of multi-sensor fusion: increased reliability and a valuable indication of the currently achieved accuracy. We provide examples for cooperative automotive applications that apply utility-based information dissemination in a vehicle-to-vehicle communications setting as well as an outlook to collaborative Simultaneous Localization and Mapping (SLAM).

1 Introduction

1.1 Principles of Multi-Sensor Fusion

Sensor fusion is based on the principle that sensors provide a means to estimate "hidden" physical processes [1]. In many cases a single sensor is sufficient to estimate a physical property, such as a temperature of a system. However, there are three reasons why one might need to draw on multiple sensors:

1. A single sensor might by itself only provide an *erroneous measurement* which does not meet the demands of the application.
2. Some aspects of the physical system might only be *observable* through the use of two or more sensors. For example, a location sensor may be insufficient to estimate the orientation of the system.
3. In order to achieve a certain level of integrity, accuracy or availability one might need to draw upon *complementary* sensors. A well known example is the fusion between inertial sensors and location sensors such as GPS.

Multi-Sensor fusion can be formulated as a static or dynamic probabilistic estimation problem. In vehicular applications we almost always face a dynamic problem where we are estimating a continuously changing random process. This problem can be formulated mathematically as a hidden Markov process and lends itself to the framework of sequential Bayesian estimation theory. Well know examples of Bayesian sequential estimators are the Kalman filter, the grid based filter, and the sequential Monte-Carlo estimator (also known as the "particle filter").

1.2 Extensions to Decision Theory

There are two main applications for sensor fusion: 1) measuring otherwise hidden properties of a physical system (e.g. tracking a vehicle's position in order to compute a toll charge); and 2) laying the foundations for some kind of decision (e.g. automatic braking).

Sometimes the boundaries between these two applications can be unclear, since any measurement of a system will lead to some kind of decision (in our example above the computed toll charge might lead to the automated debiting of an account). However, in our formulation we shall regard an application as being a decision theoretic one if the system state itself (e.g. the location of the vehicle) is likely to be a function of a prior decision (e.g. a braking manoeuvre).

Decision theory [2] incorporates the concept of *utility* to represent the relative advantage of a particular outcome resulting from a system state. For example, the outcome of a collision between two vehicles should be associated with a very high negative utility, whereas a smooth, efficient and fast transition between two road waypoints would carry a moderate positive utility. There exists a probabilistic framework between sensor fusion and decision theory (e.g. decision networks based on Bayesian networks) that computes the expected utility for all possible decisions, allowing one to choose the one with the highest expected utility. This evaluation is called deliberative decision-making since the system deliberates about the expected utility which results from a potential action in the future. The action space may, for instance, include two kinds of actions: acceleration and deceleration. Depending on the current estimation of the distance to the preceding vehicle and the relative velocity, the system evaluates the outcome of the acceleration and the deceleration, respectively. Since acceleration most probably will decrease the distance to the preceding vehicle, the resulting safety utility will decrease. On the other hand, a higher velocity will reduce the travel time and thus increase the utility for efficient movement. Finally, the calculation of the weighted average utility will show with which action a higher utility can be expected. Utility in this context might also include expected energy consumption.

This approach is very similar to human decision-making. A person who has to decide whether to take an umbrella on a walk outside (=action), will trade getting wet against unnecessarily carrying the umbrella (=utility) given the chance of rain. The latter is an estimation process which may be based simply on a view out of the window and/or an extensive study of the weather forecast.

1.3 Model mismatch

It is important to point out at this stage that sensor fusion when performed within the context of probabilistic estimation and utility-based decision theory takes into account sensor imperfections and error sources. The distinction between classical sensor fusion that is not based on probabilistic representations and that described above is that the former makes *implicit* assumptions about the nature of the errors of the sensors and any resulting computations. In contrast to this probabilistic estimation theory forces one to make *explicit* choices about the sensor error models and process models. In practice, systems are often limited by the difficulty in modelling sensor errors correctly, in particular those rare events such as sensor failure, and statistical dependencies between sensor errors (e.g. when two sensors share a common power supply or physical fixtures or are affected by similar disturbances).

In reality it is often very difficult to model the system accurately, either because the true model is unknown, or because the model is too complex to be handled by the estimation algorithms. The result is almost always a certain degree of known and unknown model mismatch, hence almost all useful estimators of non-trivial problems in the real world are *inherently suboptimal*. A particularly striking case is one where we assume measurements from a sensor to be statistically independent over time. In a static estimation context this would lead us to average the sensor readings. Consider, however, the usual situation where a GPS receiver, for example, suffers from an error which changes only slowly over time at a fixed location (such as atmospheric-induced propagation errors, or multipath errors). In this case two undesirable effects would result from averaging. Firstly, we would yield an incorrect estimate. Secondly, and just as important, our assumed probability distribution of the error would be narrower ("more certain") than warranted by the measurements. In combination, both errors can have dire effects on systems.

2 Applications to Cooperative Driving Systems

Current vehicles have dozens of built-in sensors [3]. Examples are GPS receivers, thermometer, wheel sleep sensors, and yaw rate sensors, to name only a few. In state-of-the art driver assistance systems of modern vehicles merely the sensors of that vehicle itself are exploited. This limits the sensing horizon to features of the own vehicle such as its position or its wheel slip. With radar, lidar or camera-based sensors nearby vehicles, pedestrians or obstacles can also be detected, but this is limited to entities which are located within line-of-sight.

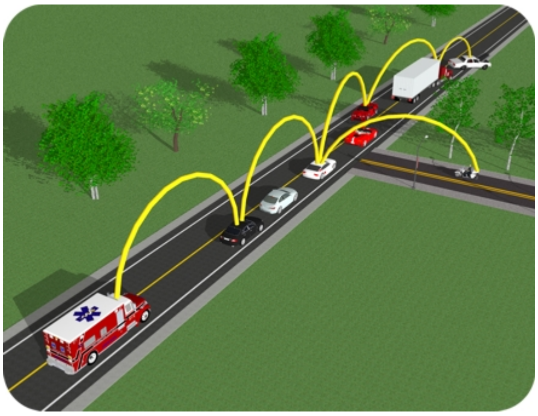

Figure 1: Vehicle-2-Vehicle Communications

With the introduction of Vehicle-2-Vehicle communications (V2V) [4], vehicles can exchange sensor measurements wirelessly and thereby significantly extend their own sensing horizon. The application of multi-hop communication even allows large-scale information dissemination up to several kilometres. Among others, multi-hop communications is essential for traffic efficiency applications such as adaptive navigation to distribute up-to-date floating car data over larger areas in order to determine optimal routes.

Exchanging sensor measurements over a wireless link between vehicles poses additional challenges to multi-sensor data fusion. In addition to the inherent measurement inaccuracy, the fusion process has to handle the loss of data, due to packet collisions, noise and interferences, and uncertainties arising from unknown sources of information. The latter can be caused by malicious intruders which try to inject false messages, such as faked traffic jams, or sensor characteristics which are a priori unknown (see section 1.3).

The key to solve these problems is the fusion of sensor measurements from remote vehicles with sensor measurements from local sensors. Since local sensors can be trusted more than remote sensors, a suitable weighting of sensor measurements in the fusion process acts as a form of consistency check. Thus, a GPS position reported by the preceding vehicle can be fused with the measurements from the local radar, lidar or camera which will improve:

1. Accuracy: Since both sensors provide independent measurements for the same hidden random process, the estimated position has a higher accuracy than a single measurement.
2. Reliability: Two independent sources of information reporting similar results can be trusted more than a single source of information.
3. Robustness: In case of malfunction or unavailability of one sensor (examples are depicted in Figure 3), the other sensor can still be used for the position estimation.

The fusion of position-related information is only one example where multi-sensor data fusion will provide an indispensable tool to future cooperative applications.

2.1 Cooperative Adaptive Cruise Control (CACC)

CACC [5][6][7][8] is one of the applications which benefits from fused sensor measurements to improve the distance estimation to the preceding vehicle. In particular in the situations depicted in Figure 3 the additional position information from the preceding vehicle is of high importance to increase robustness. But also the additional information on acceleration, driver intent and vehicle characteristics which is part of the periodic messages, also known as Cooperative Awareness Message (CAM) [9], is extremely valuable.

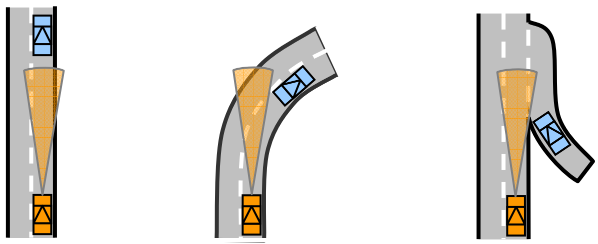

Figure 3: Radar unavailability

However, a CACC system which only focuses on the immediate preceding vehicle does not allow an application to a platoon of multiple vehicles. This fact is based on the so called shockwave effect [8] where the accumulation of delays results in an increasing collision risk the longer the platoon becomes. Since V2V communications is capable of exchanging information even outside the line-of-sight region, the vehicle can react on unforeseeable manoeuvres of vehicles in front of the immediate preceding vehicle. This approach stabilizes the platoon and, hence, increases safety, efficiency and comfort of driving.

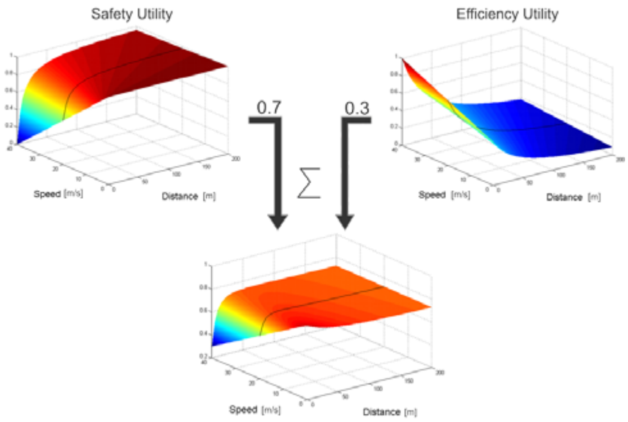

Figure 2: Weighted sum of safety and efficiency utility

Optimal decisions based on a distance estimation to the preceding vehicle with residual uncertainty require an evaluation of the action space actively taking into account this uncertainty. A further cause for uncertainty is the fact that an outcome of an action takes effect in the future. With its prediction additional uncertainty arises. With a utility-based decision-making process (see section 1.2) the current action space is evaluated by its utility in the future. Exemplary utility functions are depicted in Figure 2. From a safety perspective longer distances to the preceding vehicle have a higher utility. And, with a higher speed the distance needs to be longer in order to achieve an equal utility. In contrast to this, higher speeds are more efficient in terms of travel time because the driver reaches its destination faster. Furthermore, for efficiency a short distance is preferable. The weighted sum provides the final utility which is used in the utility-based decision-making (see section 1.2). With this approach the action selection is geared to achieve the maximum utility in the future given all past measurements.
Further utility (sub-)functions can be integrated to target multiple objectives. Examples are utility functions for energy efficiency, comfort or cooperative utility functions to optimize the behaviour of the whole cooperation (cooperative decision-making).

A valuable piece information in multi-sensor data fusion is knowledge of vehicle movements. Since vehicles normally use the road infrastructure, maps can be used to significantly reduce potential hypotheses in the position estimation. But maps have two major problems: they are not constant over time due to the construction of new roads and temporary relocation or blockages because of construction sites. Second, they are subject to inaccuracy because of measurement inaccuracies and insufficient expressivity in the discretization of the continuous course of the road. This uncertainty has to be tackled in the fusion process. On-the-fly improvements can be achieved by an online map learning based on the observed movement of the preceding vehicles.

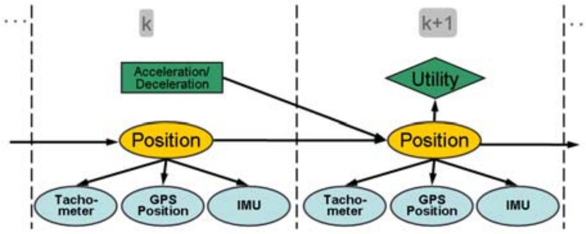

Figure 5: A simplified causal decision network. The box denotes a decision node, the diamond denotes a utility node and a function of the current state

Another example where data fusion comes into play is the reception of a black ice warning whereas the local thermometer shows a temperature of 20°C. In this case sensor fusion may come to the conclusion that the actual probability of black ice is low and a warning to the driver is suppressed. If the vehicle successively receives more and more black ice warnings from different vehicles, the actual probability of black ice increases as well as the probability that the local thermometer behaves incorrect. This kind of majority voting is nothing more than a form of multi-sensor data fusion.

2.2 Utility-based information dissemination

Each new sensor measurement has an effect on the situation estimation. With a sound modelling of the cause-effect relation between measurements (e.g. GPS position) and their corresponding situational information (e.g. real position of the vehicle) as for instance depicted in Figure 5, the worth of each new measurement can be determined [10][6]. The worth is determined by the decision, which has to be made, with the utility that is caused by the selected action. For the applications (e.g. CACC), this worth can be the trigger of an information request which is distributed via V2V communications to nearby vehicles, e.g. a request for a position measurement. With this approach each individual vehicle can determine whether a new, a priori unknown sensor measurement from a remote sensor may provide enough information to make a favourable decision. In the figurative umbrella example of section 1.2 this is equivalent to contacting a reliable weather forecast if a glance out of the window does not provide enough confidence for a decision.

A similar approach is feasible for the pro-active dissemination of sensor measurements to other vehicles via V2V communications. In case a new sensor measurement becomes available from a local sensor, the vehicle has to decide whether it want to share this information with the vehicles in the vicinity or not. Without further knowledge, the applications running on the other vehicles and, thus, also the utility functions are unknown. But, based on the information which has been distributed in the past, the vehicle can calculate the additional information included in the new measurement. The mutual information [11] can serve as a generic utility function which defines the worth of a new sensor measurement for the other vehicles. It determines how much an estimation with the additional sensor measurement differs from a pure prediction. For instance, a vehicle is stuck in a traffic jam. As long as it does not change its position, a new GPS measurement update provides little new information. But if it starts accelerating again or even changes its driving direction, a new GPS measurement has a high worth for other vehicles and, thus, shall be exchanged. A continuous exchange of all sensor measurements is disadvantageous because of the limited bandwidth which has to be shared by all vehicles in the cooperation. By calculating the worth of information for the vehicle itself as well as the vehicles which are in communication range, the communication channel is optimally utilized with an information-centric prioritization [10][6].

3 Outlook towards Simultaneous Localization and Mapping

Cartography and mapping techniques based on GPS localisation have both transformed our relationship to the physical world and have made a huge impact in the area of automotive applications. The convergence of these two technologies is supporting the uptake of specialise and mass-market location-based applications that benefit from the coupling of real-time information with maps that are more accurate and up-to-date than ever.
Many prospective location-based service (LBS) applications in ITS — including safety-critical needs for emergency — require highly accurate mapping and real-time positioning. The map enters the picture on two fronts: an accurate map improves location accuracy, and the map is needed to interpret the location in a meaningful way.

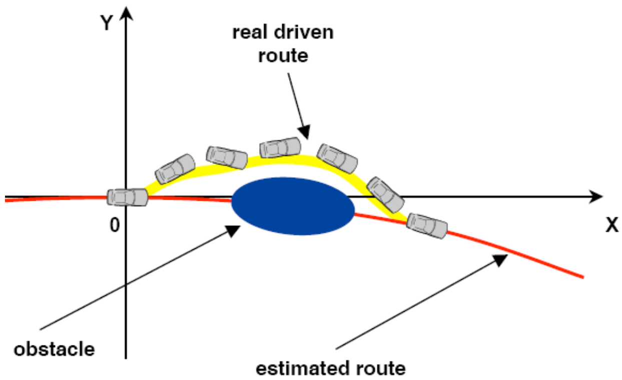

Figure 4: Obstacle Detection and Mapping (source: Straßberger [12])

In today’s rapidly changing world it cannot be taken for granted that singular organisations be responsible for maintaining maps that could in the future be updated only hours or days after changes to the environment occur. Things like temporary construction sites, pot-holes, flooded areas or other potential dangers reveal themselves through *features* which a variety of sensors (such as the human eye) can detect. Conceptually, other even more ephemeral localisable entities such as wet or icy patches or dropped goods might be mapped and their location used by all traffic participants [12] (see Figure 4). The process of mapping entities (i.e. deriving their location, shape and orientation) given the known location of several observations points is a comparatively straightforward task. Things become more difficult if the location of the sensing entity is unknown or known only to a certain degree of accuracy.

Identifying and recognising features of entities in the environment is vital in performing *Simultaneous Localization and Mapping* (SLAM) [13]. SLAM was introduced in robotics and is increasingly being used in other applications domains. The well known FastSLAM approach [14] is being used in automatic driving challenges as well as pedestrian navigation with inertial sensors [15].

One might envision a world where long-, mid- and short-term mapping are undertaken fully automatically and anonymously by vehicles whilst travelling in their environment. Formally, the process of identifying a patch of ice is “mapping” of an environmental entity. Similarly, incremental SLAM would help each vehicle to localise itself based on the current map, while at the same time performing SLAM could improve, augment and correct this map wherever necessary. Naturally, ad-hoc communication links between vehicles could be employed to process the data or coordinate the entire process. The resulting system would be less vulnerable to individual sensor, signal or signal failure (such as failure, disturbance or malicious jamming of GPS) and hence be a valuable step towards ubiquitous safety-of-life applications with critical positioning and mapping requirements.

Literature

[1] Durrant-Whyte, H., and Henderson, T. C. Multisensor Data Fusion, Springer, Handbook of Robotics. Springer Berlin Heidelberg, 2008

[2] Leonard S., The Foundations of Statistics (1954), Courier Dover Publications, 2, 1972

[3] Wallentowitz H. and Reif K.: Handbuch Kraftfahrzeugelektronik: Grundlagen, Komponenten, Systeme, Anwendungen, Vieweg+Teubner, 2006

[4] Car-2-Car Communication Consortium, http://www.car-2-car.org

[5] Röckl M., Gacnik J. and Schomerus J., Integration of Car-2-Car Communication as a Virtual Sensor in Automotive Sensor Fusion for Advanced Driver Assistance Systems, Association of German Engineers (VDI), Springer Automotive Media, International Federation of Automotive Engineering Societies (FISITA), Munich (Germany), 2008

[6] Röckl M., Cooperative Situation Awareness in Transportation, Dissertation, ISBN 978-3868535334, 2010

[7] VanderWerf J., Shladover S. and Miller M., Conceptual Development and Performance Assessment for the Deployment Staging of Advanced Vehicle Control and Safety Systems, PATH - University of California, Berkeley, 2004

[8] van Arem B., van Driel, C. J. G. and Visser, R., The Impact of Cooperative Adaptive Cruise Control on Traffic-Flow Characteristics, IEEE Transactions on Intelligent Transportation Systems, 4 [7], 2006

[9] ETSI TC ITS, http://www.etsi.org/its

[10] Röckl M., and Robertson P., Data Dissemination in Cooperative ITS from an Information-centric Perspective. IEEE International Conference on Communications (ICC), 2010 ,Cape Town (South Africa)

[11] Shannon C. and Weaver W., A Mathematical Theory of Communication, University of Illinois Press, 1949

[12] Straßberger, M. and Lasowski, R., A New Approach For Obstacle Detection Based On Dynamic Vehicle Behaviour, In Proceedings of the 10th International Forum on Advanced Microsystems for Automotive Applications, Berlin, Germany, 2006

[13] R. Smith, M. Self, and P. Cheeseman, Estimating uncertain spatial relationships in robotics. In I. J. Cox and G. T. Wilfong, editors, Autonomous robot vehicles, pages 167–193. Springer-Verlag New York, Inc., 1990

[14] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In Proceedings of the AAAI National Conference on Artificial Intelligence, Edmonton, Canada, 2002. AAAI.

[15] Robertson, Patrick and Angermann, Michael and Krach, Bernhard, Simultaneous Localization and Mapping for Pedestrians using only Foot-Mounted Inertial Sensors, in Proc. Ubicomp 2009, Orlando, Florida, USA, 30. Sep.-3. Oct. 2009.