Intelligent Information Dissemination
in Collaborative, Context-Aware Environments

Korbinian Frank*‡
*Institute of Communications and Navigation
German Aerospace Center (DLR)
Oberpfaffenhofen, Germany
korbinian.frank@dlr.de

Matthias Röckl§
§In2Soft GmbH
Munich, Germany
matthias.roeckl@In2Soft.de

Tom Pfeifer‡
‡TSSG
Waterford Institute of Technology
Waterford, Ireland
t.pfeifer@computer.org

Abstract—Today's mobile computing environments aggregate many entities, all of them with a number of internal sensors, processing applications and other, user given information that can be shared with others over the available networks. Tomorrow's ubiquitous computing environments, where the number of sensors is assumed to be even significantly higher, face the challenge that information exchange between entities has to be controlled, not only to protect privacy and to unburden the wireless networks - but also to reduce load on the receiving entities that have to process all incoming information.

The approach we propose in this work measures the importance of some information and the interest of the receiver in it, before it is transferred. We apply it in two scenarios with limited resources. In vehicle-to-vehicle communications the transmission time while cars meet and bandwidth availability of the wireless channel is the critical resource, forcing to reduce information exchange. In context inference on mobile devices processing power and battery life are limited and responsiveness to user actions is most important. Hence only the most important information should be processed.

Keywords-Bayesian Networks; Context Inference; Bayeslet; Composition; Information Dissemination; Vehicle-to-Vehicle; V2V; Car-to-Car; Vehicle ad-hoc networks; VANET; Entropy; Value of Information; Mutual Information

I. INTRODUCTION

The added value of ubiquitous computing systems stems from the fact that participating devices have access not only to their own information, but also the information of other service providers, sensors or information consumers. The information received from heterogeneous and redundant sources is usually evaluated and fused with own information to form more accurate, more reliable or new knowledge [1].

Information is exchanged via centralised architectures like the internet, but also in ad-hoc networks like for *vehicle-to-vehicle (V2V)* communications in safety-critical driver assistance systems [2] or communicating personal smart spaces [3]. In particular the ad-hoc networks in the latter situations are always wireless, may exist only for some moments and their capacity is limited. Hence, to make ubiquitous systems work, the information exchange between entities has to be reduced to the necessary minimum that does not flood and thus overburden networks and receivers.

On the other hand, to keep the advantages of the pervasively networked ubiquitous system, no necessary or useful information shall be retained. Smart algorithms have to decide thereby on sender and receiver side based on the content of the information message, what is necessary and what should not be transmitted.

For resource-constrained mobile devices also the processing overload which emerges from a plethora of information is critical. Bayesian algorithms are well-explored for the knowledge inference of such systems [4] and provide the necessary freedom in modelling situations for inference as well as the adaptability to new situations by automated learning from data. Just like for other inference approaches like logics (the satisfiability (SAT) problem is NP complete), the exact evaluation of large problems is intractable for Bayesian algorithms: the evaluation of Bayesian Networks for instance is NP hard [5].

The objective of this paper is to present a content aware information rating approach used to minimize the necessary information exchange in dynamic and resource-constrained environments. We show the applicability in V2V communications and in context inference where it reduces the network traffic and computational burden for small devices respectively.

In the following section, we will describe these two scenarios more in detail, to make obvious why an approach to reduce communications is needed. A suggested way to probabilistically fuse information from different sources is presented in section III, before section IV describes the information theory and the concrete approaches to be used in our work. Section V then applies them to a use case combining the application scenarios of section II and shows with a concrete example how they can be used. Finally, we close this paper with a short conclusion and an outlook on how we will pursue and integrate this work in the future.

II. APPLICATION SCENARIOS

A. Resource-Constrained Context Inference

Context Inference uses information coming from local and remote sensors to reason about situations or facts that

are not yet known. As mobile devices like they are used in pervasive computing are often small and have to cope with less energy and often also less processing power, inference algorithms should be able to adapt accordingly, giving correct results using a minimum of resources. A logical step in that direction is to make the algorithms work also with reduced input, selecting only the information that leads to the highest information gain and not using available information that would not add significant value or that would cause too high costs. The algorithms in section IV will provide means for that.

B. Bandwidth Limited Vehicle-to-Vehicle Communication

With the introduction of V2V communications, sensor measurements can be exchanged between vehicles, which opens up a collaborative distributed information network. Since the communication channel has to be used collaboratively by all nodes in the network, bandwidth has to be shared intelligently, in particular if safety critical applications are considered. Depending on the current situation, sensor measurements can have a high importance to a node and its neighbors (e.g. to prevent a collision). On the other hand, the transmission of information over the wireless channel consumes bandwidth and thus prevents others from transmitting their, maybe more important, information. Whether information shall be exchanged or not, hence becomes an information-theoretic problem which requires decision-making based on the information which is to be disseminated and the costs which occur due to the dissemination. The algorithms of section IV introduce concepts for adaptive information dissemination using a utility-based decision-making approach.

III. BAYESLETS FOR EFFICIENT INFORMATION FUSION FROM HETEROGENEOUS AND VOLATILE SOURCES

Bayesian networks (BN) have evolved as a major tool in a wide area of scientific disciplines requiring statistical analysis, automated reasoning or exploitation of knowledge hidden in noisy (as from erroneous sensors or uncertain data sources) or incomplete data, which can be even combined with human expert knowledge. For instance Pearl [6] and Heckerman [7] provide comprehensive introductions into theory and applications.

A particular view on BNs are *Causal Networks* [9], where dependencies represent causal influence. With this interpretation of BNs in Information Fusion, every random variable represents an information entity (e.g. context attribute, sensor or service), its values the mutually exclusive and exhaustive states of the entity. The structure and the probabilities in the conditional probability distributions (CPD) encode the existing knowledge about a certain context.

The evaluation of BNs is called *inference*, which assigns the *random variables (RV)* in question the marginalised and normalised probabilities. Inference collects the available
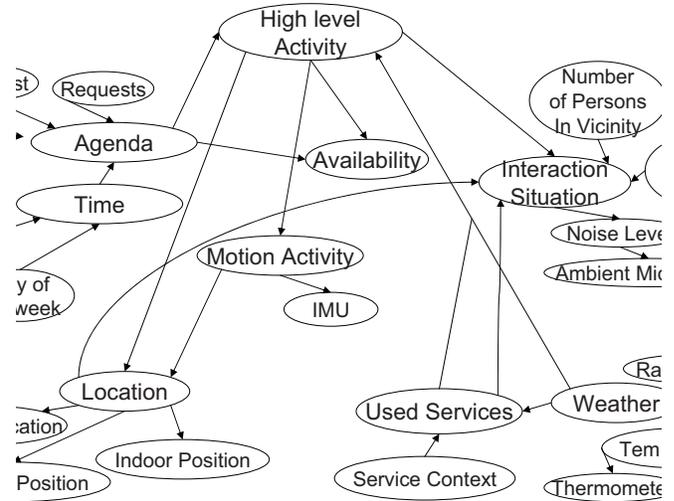

Figure 1. A cutout of the example Bayesian Network in [8] representing a view on important contextual aspects of a pervasive system user.

knowledge, adds them as *evidence* in the respective RVs and propagates the consequences in the whole network. This propagation visits every node in the graph twice, which influences the complexity of this process. To keep inference tractable although being NP hard, the factors determining the complexity, in particular the number of random variables, number of edges, number of values have to be controlled.

As the number of RVs represents the amount of included information, it makes sense to provide a mechanism that enables correct and comprehensive inference with a minimum of RVs, as then less information has to be transmitted and inference will need less computational resources.

This is achieved with *Bayeslets* [10], thematic BNs with predefined input and output nodes. They partition full BNs into thematic RV groups (see [8]) that can be managed and evaluated separately, in order to limit the number of RVs involved in one inference process. Figure 2 for instance applies this concept to the BN of Figure 1.

In scenario II-A, Bayeslets represent different inference domains, such as *Location*, *Availability* or *Environment*, or information from different users (see Figure 2). In scenario II-B Bayeslets contain information transferred from vehicles nearby, like the other vehicle's *Position* or *Direction*.

IV. UTILITY DETERMINATION

To decide whether to connect two (or more) Bayeslets, e.g. to exchange two (or more) sensor measurements between vehicles, the utility of the additional information has to be determined. In utility theory the term *utility* is defined by the value a piece of information provides to the system. Evidently, this value strongly depends on the usage of the information. In context-aware services, evidence, such as sensor measurements, is used to infer higher-level information, such as the current activity of a person or the
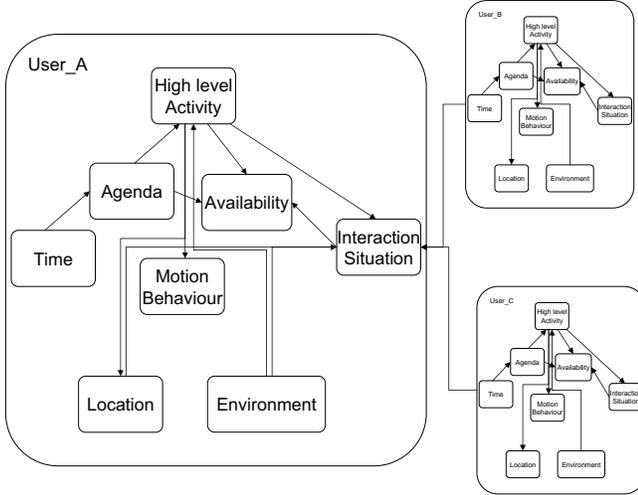
Figure 2. This is a network of Bayeslets from [8], grouping, fragmenting and extending the situation described in Figure 1. The main concepts are grouped into Bayeslets, and similar Bayeslets of relevant other persons are linked in.

probability of a rear-end collision.

The utility of an additional piece of evidence $y \in Y$ can be quantified by the difference between the utility with the additional evidence and without the additional evidence. Here, y refers to the evidence of a random variable (e.g. a measurement of an accelerometer) which is not conditionally independent of another random variable X (e.g. the activity of the person). Thus, the utility of evidence y shall be determined by its impact on X. This impact is defined by:

$$U(X : y) = \underbrace{U(X|y)}_{\substack{\text{Utility of X} \\ \text{given known } y}} - \underbrace{U(X)}_{\substack{\text{Utility of X} \\ \text{not given } y}} \quad (1)$$

with $U(X) : X \rightarrow \mathbb{R}$ as the utility function which maps random variable X to a real number. $U(X|y)$ is the utility function that maps the random variable X given evidence y to a real number. Since X is subject to uncertainty, equation 1 can be enhanced to calculate the *expected utility (EU)* gain by summing over all states weighted by their probability of occurrence:

$$EU(X : y) = EU(X|y) \quad - EU(X) \quad (2)$$
$$= \underbrace{\sum_{x \in X} U(x|y)P(x|y)}_{\substack{\text{Expected Utility of X} \\ \text{given known } y}} - \underbrace{\sum_{x \in X} U(x)P(x)}_{\substack{\text{Expected Utility of X} \\ \text{not given any } y}}$$

Unfortunately, it is not possible to determine y unless the Bayeslets are connected (in scenario II-A) or vehicles exchanged their state (in scenario II-B). The solution is to calculate the expected utility which can be gained from acquiring any $y \in Y$ instead. Therefore, one has to sum over the utility $U(x|y)$ of all possible outcomes of y, weighted

not merely by $P(x|y)$ but also by $P(y)$, the prior probability of y:

$$EU(X : Y) = EU(X|Y) - EU(X) \quad (3)$$
$$= \underbrace{\sum_{x \in X} \sum_{y \in Y} U(x|y)P(x|y)P(y)}_{\substack{\text{Expected Utility of X} \\ \text{given unknown } Y}} - \underbrace{\sum_{x \in X} U(x)P(x)}_{\substack{\text{Expected Utility of X} \\ \text{not given any } y}}$$

To account for already acquired knowledge, e.g. from already connected Bayeslets, equation 3 can be enhanced to consider contextual knowledge c:

$$EU(X : Y|c) = \quad (4)$$
$$= \underbrace{\sum_{x \in X} \sum_{y \in Y} U(x|y,c)P(x|y,c)P(y,c)}_{\substack{\text{Expected Utility of X} \\ \text{given unknown } Y \text{ and known c}}} - \underbrace{\sum_{x \in X} U(x|c)P(x|c)}_{\substack{\text{Expected Utility of X} \\ \text{given known } c}}$$

With equation 4 an intelligent decision-maker could control the exchange of information as follows:

$$\texttt{if } (EU(X : Y|c) > 0) \texttt{ use } Y \quad (5)$$
$$\texttt{else } \text{ retain } Y$$

$EU(X : Y|c) \geq 0$ as it is shown in [11]. This means that additional evidence never decreases the utility and, thus, Bayeslets in most cases shall be connected. But a connection of Bayeslets also implies drawbacks which are, for instance, higher processing load, communication costs in case Bayeslets reside on remote entities or monetary costs if evidence from a remote Bayeslet is provided by a commercial service, e.g. a commercial weather service for highly accurate weather information.

To take into account costs that reduce the gross utility, equation 4 can be enhanced to the so called *Net Expected Utility (NetEU)*:

$$NetEU(X : Y|c) = EU(X : Y|c) - C(Y) \quad (6)$$

in case costs for all $y \in Y$ are equal. If this is not the case, the $NetEU(X : Y|c)$ is calculated by:

$$NetEU(X : Y|c) = \quad (7)$$
$$= \sum_{x \in X} \sum_{y \in Y} (U(x|y,c) - C(y))P(x|y,c)P(y,c)$$
$$- \sum_{x \in X} U(x|c)P(x|c)$$

In the following we elaborate two different kinds of utility functions which can be used in the connection of Bayeslets.

A. Probability-based Utility Functions

Often the utility of random variable increases the "better it is known", i.e. the less the inherent uncertainty is. Also humans regularly acquire new information from independent sources if they are uncertain about the "true" state of an

unknown process: If we are uncertain about the weather tomorrow, we check the recent weather forecast.

A suitable utility function increases proportionally to the inverse of the certainty. The binary logarithm of the probability as used by Shannon in [12] fulfills this requirement. In this case $EU(X : Y)$ of equation 3 is equivalent to Shannon's Mutual Information $I(X : Y)$:

$$EU(X : Y) = I(X : Y) \qquad (8)$$
$$= \sum_{x \in X} \sum_{y \in Y} log_2 P(x|y) P(x|y) P(y) - \sum_{x \in X} log_2 P(x) P(x)$$
$$= -H(X|Y) + H(X) \quad \text{[in bits]}$$

with $H(X) = -E_X[\log_2 P(X)]$ being the entropy of the random variable X, $H(X|Y) = -E_{X,Y}[\log_2 P(X|Y)]$ being the conditional entropy of X given Y and $E_X[\cdot]$ being the expectation function over X.

The proof for $EU(X : Y) \geq 0$ can be found in [11].

B. Decision-based Utility Functions

Although normally it is beneficial to head towards less uncertainty in the higher-level information, in decision support systems utility shall be based on the outcome of actions instead of uncertainty reduction. Thus, we extend our model by a further feature which is called "action". Formally, Savage defines an action as *"a function attaching a consequence to each state of the world"* [13]:

$$a : X \rightarrow X \qquad (9)$$
$$a(x) = x_a, \quad \text{with } x, x_a \in X$$

with x_a being the consequence of x after action a.

A rational decision-maker will choose the action which maximizes the resulting utility. Thus, we can substitute the expected utility by the *maximum expected utility (MEU)*:

$$EU(X|Y) = MEU(X|Y) \qquad (10)$$
$$= \max_{a \in A} EU(X|Y, a) = \max_{a \in A} \sum_{x \in X} U(x) P(x|a, Y)$$

If we substitute $EU(X|Y)$ by $MEU(X|Y)$ in equation 3, $EU(X : Y)$ is equivalent to the so called *Value of Information (VoI)*:

$$EU(X : Y) = VoI(X : Y) = MEU(X|Y) - MEU(X) \qquad (11)$$

Additional evidence never decreases the VoI, as $MEU(X|Y) \geq MEU(X)$ (see the proof in [11]).

V. EXAMPLE APPLICATIONS

In the following the decision algorithm introduced in the previous sections will be shown by the exemplary application *Cooperative Adaptive Cruise Control (CACC)* [14], [15].

A. Probability-based Utility Functions

The information dissemination with a probability-based utility function (sec. IV-A) has been implemented within a sequential Monte Carlo estimator (SMC) which is also known as particle filter (more details can be found in [16], [11]). The detailed parameter specifications for the filter configuration can be found in Röckl et al. [17].

For the evaluation we implemented a time-triggered simulation environment which generates sensor output for simulated vehicles with 10Hz update rate. The evaluation scenario used in this work is a straight and a "zig-zag" road with two vehicles following each other. Each vehicle is assumed to be equipped with a GNSS receiver, an odometer and a compass which provide erroneous measurements. The preceding vehicle sends out these measurements according to the algorithms defined in section IV-A.

Fig. 3 shows the expected utility $EU(X : y)$ of the preceding vehicle during a simulation run on the "zig-zag" road. During the straight road segments the $EU(X : y)$ varies about 0.2bit around 1.1bit. Thus, there is no significant change and the message priority is low. After each sharp bend $EU(X : y)$ shows distinguishable peaks since the prediction of the future position acts on the assumption of an ongoing straight movement with an unknown map. The update with recent evidence provides an unexpected change in this movement and thus has a high $EU(X : y)$. The peak height depends on the current measurement quality which is not shown in the figure.
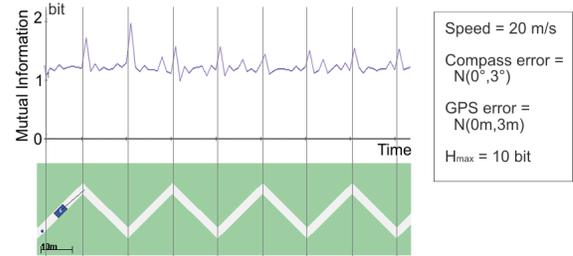

Figure 3.   Mutual information on a "zig-zag" road with sharp 90° bends

Fig. 4 shows the progress of the $EU(X : y)$ for the target vehicle driving on a straight road. Due to faults in the positioning system, GNSS updates arrive only with an update rate of 0.4Hz. Every time a new measurement becomes available a peak in $EU(X : y)$ can be recognized since the uncertainty of the prediction is significantly reduced by the new measurement.

Up to now only the variation of $EU(X : y)$ depending on the measurement value and its uncertainty has been inspected. In fig. 5 the actual resulting update rate is evaluated when the $EU(X : y)$ is used to decide whether to disseminate the measurement. If the measurement was retained according to the else branch of eq. 5, its $EU(X : y)$
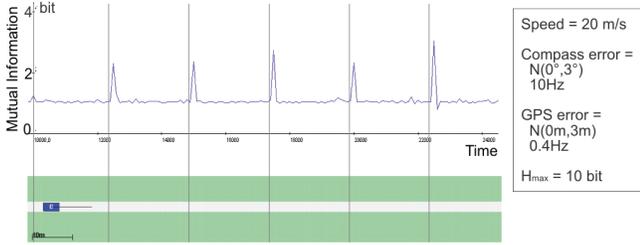
Figure 4. Mutual information on a straight road with GNSS position fixes every 2.5 seconds

value was added to the expected utility of the upcoming measurement resulting in a cumulative expected utility value. The costs for the dissemination were uniformly set to 10. Thus, every time the cumulated $EU(X : y) > 10$ the message containing the evidence position, velocity and heading is disseminated.

The continuous line depicts the velocity of the preceding vehicle. The vertical bars show the mean update rate during the acceleration, the constant high speed, the deceleration and the standstill phase. During the standstill the update rate is the lowest with ca. 0.4Hz. During the acceleration and deceleration phase the update rate raises to 0.9-1Hz and during the high speed movement the update rate reaches its maximum with ca. 1.1Hz. This is justified due to the higher uncertainty of the movement prediction with higher speeds.
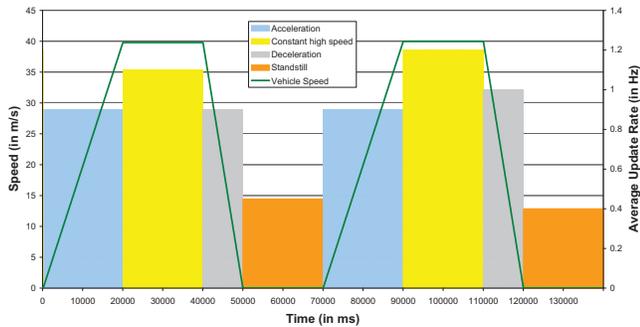

Figure 5. Average update rate during the driving maneuvers *acceleration*, *constant high speed*, *deceleration* and *standstill*

B. Decision-based Utility Functions

On a road with multiple lanes in the same direction the ego vehicle has to determine the actual lane the preceding vehicle is driving on. If, for instance, the ego vehicle is located on the center lane of a three-laned road, the preceding vehicle can be located on the same, the right or the left lane. These three states are denoted as *center*, *right* and *left* of the Lateral Distance (LD) in the following. Evidence to determine the preceding vehicle's position is given by Radar Lateral Measurements (RLM) and V2V

Lateral Measurements (VLM) (with GNSS information transmitted via V2V communications) with the same states as LD. Additionally, an overall utility function, which represents a weighted average of the Safety and Efficiency utility with a 3:1 weighting is specified. The decision differentiates the two states *accelerate* and *decelerate*. The whole probabilistic decision network is depicted in Figure 6.
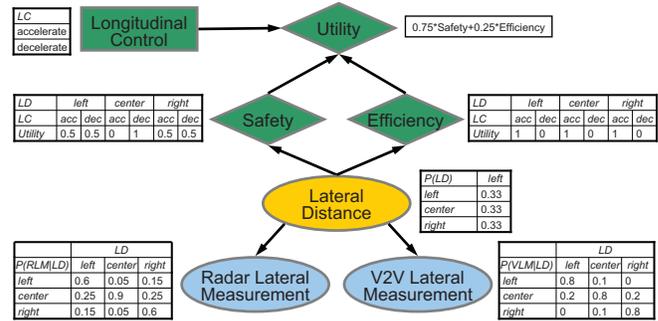

Figure 6. *Longitudinal Control* decision based on the *Lateral Distance* with evidence from radar and V2V

If no evidence is available, the state of LD is uniformly distributed, each state with a probability of 1/3 (see left-most depiction in figure 7). In this case there is a tie between the actions *accelerate* and *decelerate*. Both have a utility of 0.5. In order to decide whether another Bayeslet, including additional sensor measurements either from the radar (RLM) or V2V (VLM), shall be connected the VoI can be calculated, which provides a value greater than zero to both evidences. A RLM has a value of 0.11 and thus carries more information than a VLM of 0.10. These calculations are independent of the unknown actual state of the evidence. Thus, the decision algorithm requests a radar measurement (e.g. by sending the according CAN RTR frame on the respective CAN bus) or listening for an appropriate CAN frame on the respective bus.

When the radar measurement is received, the input node RLM can be updated with the new evidence. If, for instance, the radar measured a *center* state, the VoI of the VLM reduces to 0.07 because both evidences are dependent due to the common cause LD. Thus, the belief in the *center* state as the actual outcome of the VLM already has a probability of 59%. But the VoI of VLM is still positive and, hence, is expected to provide a value for the decision-making. This is justified since after the incorporation of the *center* state from RLM and an additional acquisition of VLM which may provide a different state than *center* (a *left* state in the right-most depiction of Figure 7) the MEU switches over to 0.51 for the *accelerate* action. Thus, the best action changes from *decelerate* to *accelerate* by the acquisition of the VLM. Evidently, in this example requesting position information from the preceding vehicle via V2V communications may provide
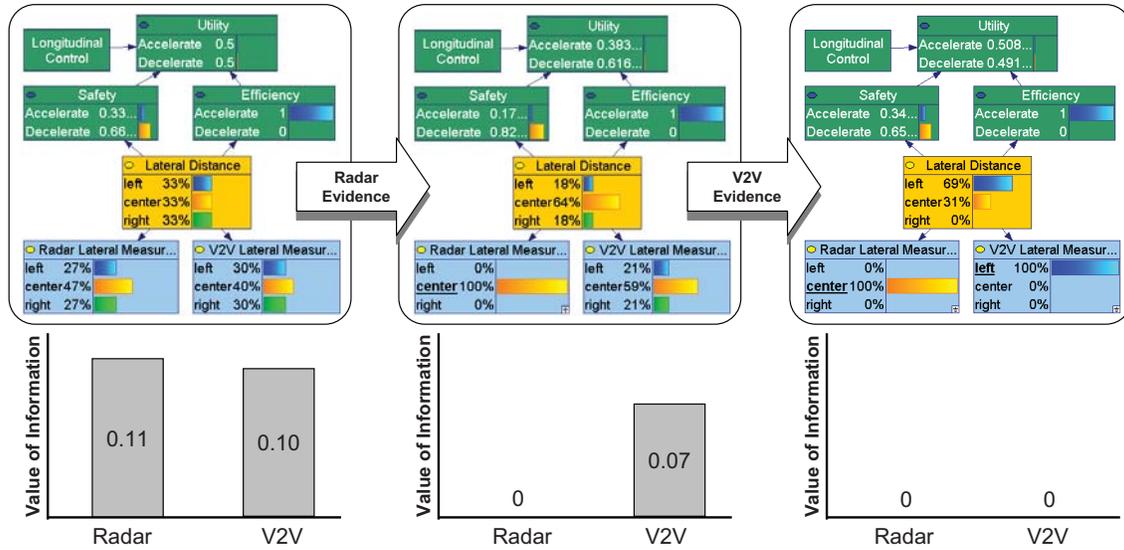
291

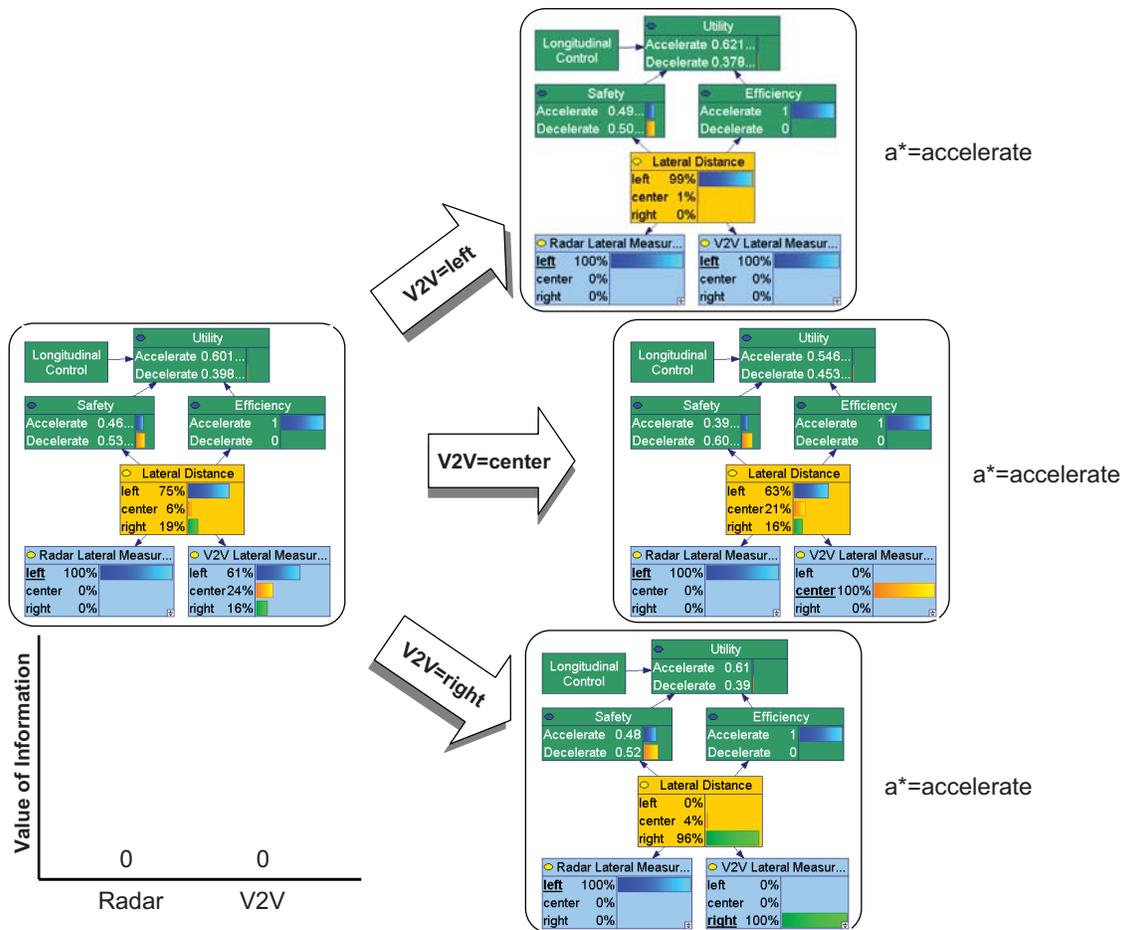Figure 7. Variation of the Value of Information with evidences from radar and V2V communications

Figure 8. Variation of the maximum expected utility (MEU) with different outcomes of the V2V evidence after observing the radar evidence *left*

292

a valuable benefit for the decision-making (neglecting any transmission costs).

Thereby VoI is not a general evaluation on a per-sensor basis, but is based on the actual outcome level as explained in the previous section. This can be seen in Figure 8. The figure is based on the same parameters as the previous evaluation, but instead of the *center* state as the outcome of the RLM the radar provided a *left* measurement. The best action a^* accordingly is *accelerate* with a MEU of 0.60. In this case the VoI for the VLM is 0 whereas in the previous analysis it was 0.07. The figure also shows the situation after the acquisition of VLM . If it provides as well the *left* state the best action evidently is *accelerate*. If it provides the *center* state still *accelerate* is the best action and even if it provides a *right* measurement the best action is *accelerate*. Thus, independent of the outcome of the VLM the best action is always *accelerate*, the VoI is zero and hence this Bayeslet is not requested, saving transmission time and bandwidth.

VI. Conclusion and Future Work

This work shows that the theory of intelligent, content-based information dissemination works in a real situation as demonstrated with the both approaches using decision-based and probability-based utility functions. These approaches can save both bandwidth and computation time, as only the necessary information is considered.

In the future this work will have to be tested not only in theory, but in realistic traffic simulations and reality, as well as in different use cases.

Acknowledgments

References

[1] T. Pfeifer, "Redundant positioning architecture," *Computer Communications, Elsevier Science Publishers B.V. North-Holland*, vol. 28, no. 13, pp. 1575–1585, August 2005.

[2] M. Röckl, K. Frank, T. Strang, M. Kranz, J. Gacnik, and J. Schomerus, "Hybrid fusion approach combining autonomous and cooperative detection and ranging methods for situation-aware driver assistance systems," in *Proceedings of International Symposium on PIMRC 2008*. IEEE, September 2008. [Online]. Available: http://elib.dlr.de/58920/

[3] M. Crotty, N. Taylor, H. Williams, K. Frank, I. Roussaki, and M. Roddy, "A pervasive environment based on personal self-improving smart spaces," in *Architectures and platforms for AMI, workshop on AmI08, European conference on Ambient Intelligence*, November 2008. [Online]. Available: http://elib.dlr.de/57797/

[4] D. Wagner and R. Wattenhofer, Eds., *Algorithms for Sensor and Ad Hoc Networks, Advanced Lectures*, ser. Lecture Notes in Computer Science, vol. 4621. Springer, 2007.

[5] G. F. Cooper, "Probabilistic inference using belief networks is NP-hard," Medical Computer Science Group, Knowledge Systems Laboratory, Stanford University, Stanford, CA, Tech. Rep. KSL-87-27, 1987.

[6] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco: Morgan Kaufmann, 1988.

[7] D. Heckerman, "A Tutorial on Learning with Bayesian Networks," Microsoft Research, Redmond, Washington, Technical Report MSR-TR-95-06, 1995. Revised June 1996. [Online]. Available: http://citeseer.ist.psu.edu/41127.html

[8] K. Frank, M. Röckl, and T. Pfeifer, "Optimizing dynamic composition of bayesian networks for context sensing and inference," in *the 35th Annual IEEE Conference on Local Computer Networks (LCN)*, Denver, Colorado, USA, 2010.

[9] J. Pearl, *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.

[10] K. Frank, M. Röckl, and P. Robertson, "The Bayeslet concept for modular context inference," in *Proceedings of Second International Conference on UBICOMM08*, J. Lloret Mauri et al., Ed. IEEE CPS, May 2008, pp. 96–101. [Online]. Available: http://elib.dlr.de/54515/

[11] M. Röckl, "Cooperative situation awareness in transportation," Ph.D. dissertation, University of Innsbruck (Austria), 2010.

[12] C. E. Shannon and W. Warren, "A Mathematical Theory of Communication," 1949.

[13] L. Savage, *The Foundations of Statistics (1954)*, 2nd ed. Courier Dover Publications, 1972.

[14] M. Röckl, J. Gacnik, and J. Schomerus, "Integration of Car-2-Car Communication as a Virtual Sensor in Automotive Sensor Fusion for Advanced Driver Assistance Systems," in *FISITA08*. Munich, Germany: Springer, 2008.

[15] B. van Arem, C. J. G. van Driel, and R. Visser, "The Impact of Cooperative Adaptive Cruise Control on Traffic-Flow Characteristics," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 4, 2006.

[16] M. Röckl, K. Frank, T. Strang, M. Kranz, J. Gacnik, and J. Schomerus, "Hybrid fusion approach combining autonomous and cooperative detection and ranging methods for situation-aware driver assistance systems," in *IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC) 2008*, September 2008.

[17] M. Röckl and P. Robertson, "Data dissemination in cooperative its from an information-centric perspective," in *IEEE International Conference on Communications (ICC)*, Cape Town, South Africa, 2010.