

Inter-Vehicular Communication Towards Cooperative Driving

Edited by

Onur Altintas¹, Suman Banerjee², Falko Dressler³, and
Geert Heijenk⁴

1 TOYOTA InfoTechnology Center USA – Mountain V, US,
onur@us.toyota-itc.com

2 University of Wisconsin – Madison, US, suman@cs.wisc.edu

3 Universität Paderborn, DE, dressler@ccs-labs.org

4 University of Twente, NL, geert.heijenk@utwente.nl

Abstract

Looking back at the last decade, one can observe enormous progress in the domain of vehicular networking. In this growing community, many ongoing activities focus on the design of communication protocols to support safety applications, intelligent navigation, multi-player gaming and others. This seminar shifted the focus from basic networking principles to networked control applications. We were particularly interested in eSafety applications and traffic efficiency applications that are thought to yield substantial benefits for the emerging “cooperative automated driving” domain. The seminar brought together experts from several fields, including classical computer science (computer networking, simulation and modeling, operating system design), electrical engineering (digital signal processing, communication networks), and automated driving (mechanical engineering, image processing, control theory), to discuss the most challenging issues related to inter-vehicular communication and cooperative driving.

Seminar May 13–16, 2018 – <http://www.dagstuhl.de/18202>

2012 ACM Subject Classification Networks → Cyber-physical networks, Networks → Mobile networks, Networks → Network architectures, Networks → Network performance evaluation, Networks → Network protocols

Keywords and phrases automated driving, cooperative driving, road traffic safety, vehicular networking

Digital Object Identifier 10.4230/DagRep.8.5.31

Edited in cooperation with Isabel Wagner

1 Executive Summary

Falko Dressler (Universität Paderborn, DE)

Onur Altintas (TOYOTA InfoTechnology Center USA – Mountain V, US)

Suman Banerjee (University of Wisconsin – Madison, US)

Geert Heijenk (University of Twente, NL)

Katrin Sjöberg (Volvo, Sweden)

License © Creative Commons BY 3.0 Unported license

© Falko Dressler, Onur Altintas, Suman Banerjee, Geert Heijenk, and Katrin Sjöberg

Looking back at the last decade, one can observe enormous progress in the domain of vehicular networking. In this growing community, many ongoing activities focus on the design of



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 3.0 Unported license

Inter-Vehicular Communication Towards Cooperative Driving, *Dagstuhl Reports*, Vol. 8, Issue 05, pp. 31–59

Editors: Onur Altintas, Suman Banerjee, Falko Dressler, and Geert Heijenk



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

communication protocols to support safety applications, intelligent navigation, multi-player gaming and others. Very large projects have been initiated to validate the theoretic work in field tests and protocols are being standardized. With the increasing interest from industry, security and privacy have also become crucial aspects in the stage of protocol design in order to support a smooth and carefully planned roll-out. We are now entering an era that might change the game in road traffic management. This is supported by the U.S. federal government announcement in December 2016 that National Highway Traffic Safety Administration (NHTSA) plans to make V2V devices in new vehicles mandatory. This coincides with the final standardization of higher layer networking protocols in Europe by the ETSI.

The vehicular networking research also complements the ongoing activities towards automated driving. Very successful activities started with the Google and lead to first projects on the road such as the Singapore driverless taxi service or the platooning experiments in Scandinavia and now Germany.

The management and control of network connections among vehicles and between vehicles and an existing network infrastructure is currently one of the most challenging research fields in the networking domain. Using the terms Vehicular Ad-hoc Networks (VANETs), Inter-Vehicle Communication (IVC), Car-2-X (C2X), or Vehicle-2-X (V2X), many applications – as interesting as challenging – have been envisioned and (at least) partially realized. In this context, a very active research field has developed. There is a long list of desirable applications that can be grouped into four IVC categories:

1. eSafety applications that try to make driving safer, e.g. road hazard warning;
2. traffic efficiency applications aiming at more efficient and thus greener traffic, e.g., detection of traffic jams;
3. manufacturer oriented applications, e.g., automatic software updates; and
4. comfort applications, e.g. automatic map updates.

In 2010, a first Dagstuhl Seminar (10402) was organized on the topic of inter-vehicular communication. The motivation was to bring together experts in this field to investigate the state of the art and to highlight where sufficient solutions already existed. The main outcome of this very inspiring seminar was that there are indeed areas within this research where scientific findings are being consolidated and adopted by industry. This was the consensus of quite intriguing discussions among participants from both industry and academia. Yet, even more aspects have been identified where substantial research is still needed. These challenges have been summarized in the following IEEE Communications Magazine article [1].

A follow-up seminar (13392) was organized in 2013. The goal was to again bring together leading researchers both from academia and industry to discuss if and where the previously identified challenges have been adequately addressed, and to highlight where sufficient solutions exist today, where better alternatives need to be found, and also to give directions where to look for such alternatives. Furthermore, it was the goal of this workshop to go one step beyond and identify where IVC can contribute to the basic foundations of computer science or where previously unconsidered foundations can contribute to IVC. It turned out that quite a number of research questions were still open or insufficiently addressed. This particularly included scalability and real-time capabilities. These challenges have been summarized in the following IEEE Communications Magazine article [2].

We now shifted the focus of this seminar from basic networking principles to networked control applications. We were particularly interested in the first two IVC categories that are thought to yield substantial benefits for the emerging “cooperative automated driving” domain. It is of utmost importance to bring together expertise from classical computer

science (computer networking, simulation and modeling, operating system design), from electrical engineering (digital signal processing, communication networks), as well as from automated driving (mechanical engineering, image processing, control theory). Building upon the great success of the first two seminars, with this follow-up seminar, we aimed to again bring together experts from all these fields from both academia and industry.

The seminar focused intensively on discussions in several working groups. To kick-off these discussions, we invited two keynote talks “Cooperative Driving A Control of a Networking Problem?” by Renato Lo Cigno and “Cooperative driving – maneuvers, perception, and IVC” by Lars Wolf. These keynotes were complemented by four additional talks: Human-in-the-Loop: Towards Deeply Integrated Hybridized Systems (Falko Dressler), Machine Learning for Cooperative Driving (Geert Heijenk), Measuring Privacy in Vehicular Networks (Isabel Wagner), and Predictable V2X Networking for Application-Networking Co-Design (Hongwei Zhang). We finally organized the following working groups on some of the most challenging issues related to inter-vehicular communication and cooperative driving:

- Ultra-Reliable Low-Latency and Heterogeneous V2X Networking,
- Human-in-the-Loop,
- Safety-critical Vehicular Network Applications,
- Security and Privacy,
- Network and Cloud based Control, and
- Sensing and Data Management.

For most of these working groups, we provide in-depth feedback from the experts in this report.

References

- 1 Falko Dressler, Hannes Hartenstein, Onur Altintas, and Ozan K. Tonguz. Inter-Vehicle Communication – Quo Vadis. *IEEE Communications Magazine*, 52(6):170–177, June 2014.
- 2 Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz, and Lars Wischhof. Research Challenges in Inter-Vehicular Communication – Lessons of the 2010 Dagstuhl Seminar. *IEEE Communications Magazine*, 49(5):158–164, May 2011.

2 Table of Contents

Executive Summary

Falko Dressler, Onur Altintas, Suman Banerjee, Geert Heijenk, and Katrin Sjöberg 31

Overview of Talks

Human-in-the-Loop: Towards Deeply Integrated Hybridized Systems <i>Falko Dressler</i>	36
Machine Learning for Cooperative Driving <i>Geert Heijenk</i>	36
Cooperative Driving A Control of a Networking Problem? <i>Renato Lo Cigno</i>	37
Measuring Privacy in Vehicular Networks <i>Isabel Wagner</i>	37
Cooperative driving – maneuvers, perception, and IVC <i>Lars Wolf</i>	38
Predictable V2X Networking for Application-Networking Co-Design <i>Hongwei Zhang</i>	39

Working groups

Ultra-Reliable Low-Latency (URLL) and Heterogeneous V2X Networking <i>Onur Altintas, Ali Balador, Suman Banerjee, Claudia Campolo, Sinem Coleri Ergen, Eylem Ekici, Sonia Heemstra de Groot, Thorsten Hehn, Florian Klingler, Renato Lo Cigno, Jörg Ott, Elmar Schoch, Jonathan Sprinkle, Erik Ström, Lars Wischhof, Andrea Zanella, and Hongwei Zhang</i>	40
Sensing and Data Management <i>Suman Banerjee, Aruna Balasubramanian, Sonia Heemstra de Groot, Albert Held, Frank Kargl, Renato Lo Cigno, Thomas Strang, Lars Wolf, and Andrea Zanella</i>	47
New Use Cases <i>Sinem Coleri Ergen, Onur Altintas, Ali Balador, Suman Banerjee, Claudia Campolo, Falko Dressler, Eylem Ekici, Sonia Heemstra de Groot, Geert Heijenk, Renato Lo Cigno, Michele Segata, Christoph Sommer, Jonathan Sprinkle, Andrea Zanella, and Hongwei Zhang</i>	49
Human-in-the-Loop <i>Falko Dressler, Eylem Ekici, Thorsten Hehn, Renato Lo Cigno, Christoph Sommer, and Lars Wischhof</i>	50
Safety-critical Vehicular Network Applications <i>Geert Heijenk, Michele Segata, Christoph Sommer, Thomas Strang, Lars Wolf, and Andrea Zanella</i>	52
Security and Privacy <i>Frank Kargl, Albert Held, Elmar Schoch, Christoph Sommer, Thomas Strang, Isabel Wagner, and Andrea Zanella</i>	54

Simulation, Modeling, and Testing

Christoph Sommer, Wai Chen, Geert Heijenk, Michele Segata, Jonathan Sprinkle,
Erik Ström, Isabel Wagner, and Hongwei Zhang 57

Participants 59

3 Overview of Talks

3.1 Human-in-the-Loop: Towards Deeply Integrated Hybridized Systems

Falko Dressler (Universität Paderborn, DE)

License © Creative Commons BY 3.0 Unported license
© Falko Dressler

Main reference Falko Dressler: “Cyber Physical Social Systems: Towards Deeply Integrated Hybridized Systems”, in Proc. of the 2018 International Conference on Computing, Networking and Communications, ICNC 2018, Maui, HI, USA, March 5-8, 2018, pp. 420–424, IEEE Computer Society, 2018.

URL <http://dx.doi.org/10.1109/ICCNC.2018.8390404>

This talk is about issues raising up when considering not 100% optimal technical systems optimized for both individual behavior and global metrics but also considering the impact of the human-in-the-loop. Technically, we are observing a paradigm shift from classical Cyber Physical Systems (CPS) to Cyber Physical Social Systems (CPSS). Humans impact our technical systems, here we talk about (semi-)automated cooperative driving, in quite many dimensions. This includes the driving behavior that depends on the driver’s demands or wishes, experiences, and capabilities that also vary over time. This is complemented by the often-cited incapability of humans to self-assess their abilities. A final frontier might be public acceptance on a global level, which might push or kill (optimal) technical solutions.

3.2 Machine Learning for Cooperative Driving

Geert Heijenk (University of Twente, NL)

License © Creative Commons BY 3.0 Unported license
© Geert Heijenk

Joint work of Geert Heijenk, Aashik Chandramohan

Machine learning is currently being introduced for self-driving cars. To a large extent the machine learning is used to interpret sensor information, including radar, lidar and video. In our research we are exploring to what extent machine learning can be used for vehicles to perform automated cooperative driving / maneuvering using information obtained through V2X communications. To this end, we need an environment for training systems and for testing systems. We therefore use a traffic simulation environment, SUMO, in which we can control the maneuvering of one, several, or all cars using machine learning agents.

In current experiments, we are using deep Q-learning to control the maneuvers of a vehicle on a 2-lane highway, where other cars are driven using one of the traditional SUMO driving models. The state input to the machine learning agent consists of the velocity and lane of the ego vehicle, and lane, and velocity of all the surrounding vehicles, assuming this is communicated using V2X. The actions the machine learning agent can take are lane change, acceleration and deceleration. The most critical part is the reward system. Currently, we give a strong negative reward for a collision. We also give negative rewards for near-collisions, and for violating traffic rules. Positive rewards are given depending on the speed, below the speed limit. In training periods of thousands of episodes, we can see the collision rate decreasing with the length of the training period, but not yet to an extent that we achieve reasonably low collision rates.

We plan to look into multi-agent learning as a way to improve the performance. Further, we are exploring other scenarios, such as intersection traffic. In that scenario, it is interesting

to see what happens if we extend the case where one car is driven by a machine learning agent and the others obey traffic rules, to the situation where all cars are driven by machine learning. From that situation maybe a new set of traffic rules will automatically emerge.

Overall, we are interested to see what is achievable using machine learning for cooperative driving, and what is the influence of all the design and parameter choices in machine learning on the learning outcome. Furthermore, we are interested to assess the potential of a simulated traffic environment for learning and for testing the outcome of the learning algorithms.

3.3 Cooperative Driving A Control of a Networking Problem?

Renato Lo Cigno (University of Trento, IT)

License © Creative Commons BY 3.0 Unported license
© Renato Lo Cigno

Joint work of Renato Lo Cigno, Michele Segata

This short talk wants to rise the attention on the fact that cooperative driving is a very complex, multi-disciplinary topic that requires a holistic approach to find a solution. All too often researchers from a specific discipline see cooperative driving shrinking the focus to their discipline, Control, Networking, Consensus, Automotive, ... losing the big picture, and also doing modeling simplifications to tackle the problem that indeed introduce biases and errors, leading to partial, if not fully wrong, solutions.

It is clear that we are missing theoretical models that are able to grab the complexity of this system, and this aspect requires attention otherwise we risk doing research that is doomed to irrelevance, because solutions will find their way into life through other means.

Another topic of attention and interest is the lack of a sort of “standardization” at the coordination level. While networking and communications are used to have standards that define the minimum set of capabilities required to enable interaction and cooperation, it seems that at the consensus, control and coordination level there is nothing like this, so that we risk to have plenty of potential solutions from different automakers that, although they are formally compatible as they use the same networking and information exchange layer, they are not actually compatible, as they apply different logics and algorithms that lead to sub-optimal decisions, or even to contrasting decisions that may even lead to dangerous situations.

3.4 Measuring Privacy in Vehicular Networks

Isabel Wagner (De Montfort University – Leicester, GB)

License © Creative Commons BY 3.0 Unported license
© Isabel Wagner

Joint work of Isabel Wagner, Yuchen Zhao

Main reference Yuchen Zhao and Isabel Wagner: “On the Strength of Privacy Metrics for Vehicular Communication”, IEEE Transactions on Mobile Computing, 2018.

URL <https://doi.org/10.1109/TMC.2018.2830359>

Vehicular communication plays a key role in near-future automotive transport, promising features such as increased traffic safety and wireless software updates. However, vehicular communication can expose drivers’ locations and thus poses privacy risks. Many schemes have been proposed to protect privacy in vehicular communication, and their effectiveness is usually

evaluated with privacy metrics. However, different privacy metrics have not been compared to each other, and it is unknown how strong the metrics are. In this talk, I evaluate and compare the strength of 41 privacy metrics in terms of four novel criteria: Privacy metrics should be monotonic, i.e., indicate decreasing privacy for increasing adversary strength; their values should be spread evenly over a large value range to support within-scenario comparability; and they should share a large portion of their value range between traffic conditions to support between-scenario comparability. I evaluate all four criteria on real and synthetic traffic with state-of-the-art adversary models and create a ranking of privacy metrics. The results indicate that no single metric dominates across all criteria and traffic conditions. I therefore recommend to use metrics suites, i.e., combinations of privacy metrics, when evaluating new privacy-enhancing technologies.

3.5 Cooperative driving – maneuvers, perception, and IVC

Lars Wolf (TU Braunschweig, DE)

License  Creative Commons BY 3.0 Unported license

© Lars Wolf

Joint work of Lars Wolf, Hendrik-Jörn Günther, Bernd Lehmann

Inter-Vehicular communication (IVC) can enable manifold types of cooperation between traffic participants, including human-driven vehicles, future autonomous vehicles, and also others like pedestrians and bicyclists. In daily live, humans cooperate and help each other in various ways, sometimes due to altruistic reasons or hoping for (indirect) reciprocity. This leads to many questions like: Can vehicular networks support such cooperation? What are the requirements for that and which new techniques are needed? Are specific methods for trust and reputation necessary? Vehicles may help others by cooperative sensing – how can such an architecture look like. Do autonomous vehicles, where no human assesses data, lead to additional demands?

Cooperative driving needs information about (i) the current situation consisting of the own perception / sensing as well as of collective perception / sensing (ii) intention of others, i.e., currently planned trajectories as well as potentially desired trajectories.

For (i) collective perception, several questions have to be solved such as: Which observations should be transmitted? At which granularity / which detail level (sensor data, objects, ...)? How and how often should transmission take place? How much does it improve the awareness ratio? How about reliability, trustworthiness, ...?

Information about (ii) intention of others enables maneuver coordination and, hence, extended cooperation. A general framework, supporting different kinds of scenarios, should be provided; thus, not for a specific traffic situation only. This requires the exchange of behavior composed of two components: a) the currently planned trajectory and b) a desired trajectory, representing a favored trajectory of a vehicle in case the need to deviate from the currently planned trajectory is detected. While maneuver coordination can be very helpful, it opens up many questions, e.g., regarding potential ambiguities, maneuver cascading and oscillation, complexity, and reliability. And there various IVC research concerns which need further study, e.g.: How to enable Maneuver Coordination Message exchange? What are the communication requirements? Which communication technologies should be used? How to deal with the interrelation between coordination necessity for increasing traffic density?

3.6 Predictable V2X Networking for Application-Networking Co-Design

Hongwei Zhang (Iowa State University, US)

License  Creative Commons BY 3.0 Unported license
© Hongwei Zhang

V2X communication is a basic enabler of the Connected-and-Automated-Vehicle (CAV) vision. In supporting safety-critical applications yet subject to complex dynamics and uncertainties, it is important to ensure predictable V2X communication (e.g., in reliability, timeliness, and throughput) so that predictable and trustworthy CAV systems can be developed. In this talk, I will present an integrated architecture for CAV applications and networks, and I will present field-deployable approaches to ensuring predictable communication reliability, timeliness, and throughput in highly-dynamic V2X networks. I will also present the applications of our architecture and algorithms to networked AR and networked control for CAVs.

References

- 1 Chuan Li, Hongwei Zhang, Jayanthi Rao, Le Yi Wang, George Yin, Cyber-Physical Scheduling for Predictable Reliability of Inter-Vehicle Communications, short paper, ACM/IEEE International Conference on Internet-of-Things Design and Implementation (IoTDI), 2018
- 2 Yu Chen, Hongwei Zhang, Nathan Fisher, Le Yi Wang, George Yin, Probabilistic Per-Packet Real-Time Guarantees for Wireless Networked Sensing and Control, IEEE Transactions on Industrial Informatics, 14(5):2133–2145, 2018
- 3 Hongwei Zhang, Xiaohui Liu, Chuan Li, Yu Chen, Xin Che, Feng Lin, Le Yi Wang, George Yin, Scheduling with Predictable Link Reliability for Wireless Networked Control, IEEE Transactions on Wireless Communications (TWC), 16(9):6135–6150, 2017
- 4 Hongwei Zhang, Xin Che, Xiaohui Liu, Xi Ju, Adaptive Instantiation of the Protocol Interference Model in Wireless Networked Sensing and Control, ACM Transactions on Sensor Networks (TOSN), 10(2), January 2014

4 Working groups

4.1 Ultra-Reliable Low-Latency (URLL) and Heterogeneous V2X Networking

Onur Altintas (TOYOTA InfoTechnology Center USA – Mountain V, US), Ali Balador (RISE SICS – Västerås, SE), Suman Banerjee (University of Wisconsin – Madison, US), Claudia Campolo (University Mediterranea of Reggio Calabria, IT), Sinem Coleri Ergen (Koc University – Istanbul, TR), Eylem Ekici (Ohio State University – Columbus, US), Sonia Heemstra de Groot (TU Eindhoven, NL), Thorsten Hehn (Volkswagen AG – Wolfsburg, DE), Florian Klingler (Universität Paderborn, DE), Renato Lo Cigno (University of Trento, IT), Jörg Ott (TU München, DE), Elmar Schoch (BMW AG – München, DE), Jonathan Sprinkle (NSF – Alexandria, US), Erik Ström (Chalmers University of Technology – Göteborg, SE), Lars Wischhof (Hochschule München, DE), Andrea Zanella (University of Padova, IT), and Hongwei Zhang (Iowa State University, US)

License © Creative Commons BY 3.0 Unported license

© Onur Altintas, Ali Balador, Suman Banerjee, Claudia Campolo, Sinem Coleri Ergen, Eylem Ekici, Sonia Heemstra de Groot, Thorsten Hehn, Florian Klingler, Renato Lo Cigno, Jörg Ott, Elmar Schoch, Jonathan Sprinkle, Erik Ström, Lars Wischhof, Andrea Zanella, and Hongwei Zhang

4.1.1 CAV Applications and Co-Design with URLL V2X Networking

Vehicle-to-Everything (V2X) communication is a basic enabler to support applications for connected and automated vehicles (CAVs). CAV applications highly vary in terms of delivery requirements, e.g., reliability, timeliness, and throughput [1, 2, 3, 4]. Among them, many high-impact applications can benefit from predictable ultra-reliable and low-latency (URLL) V2X networking. The following are some examples:

- Network/cloud-assisted control of vehicles;
- Collaborative simultaneous-localization-and-mapping (SLAM) across traffic infrastructures and vehicles;
- Human-in-the-loop Augmented Reality (AR)-assisted driving;
- Real-time CAV control such as distributed collision avoidance and cooperative adaptive cruise control.

Achieving high reliability and low-latency is typically hindered by the fundamental trade-off between such metrics [5]. They can have different impact on CAV applications, and it is difficult to know what individual CAV applications require exactly at design time in general. Therefore, it is important to enable on-the-fly characterization of communication metrics and the control/choice of reliability-timeliness trade-off by applications.

The probabilistic nature of wireless communications and inherent uncertainties in V2X networks shall be considered in application and networking co-design. For instance, the interface between applications (e.g., CAV control) and networking shall capture the nature of random communication/networking topology, and this shall be differentiated from potentially random CAV coordination/control topologies, too. To better utilize V2X communication resources, it is also important to decide whether raw data or processed/fused data are to be exchanged between CAVs.

4.1.2 URLL V2X Networking: Models and Approaches

To support application-networking co-design and in capturing the probabilistic nature of wireless communication, it is crucial to model performance metrics. One approach is as follows:

$$\text{Probability}\{D \leq D_0\} \geq P_0, \quad (1)$$

where D is the communication delay from one CAV to another (via short-range single-/multi-hop links, and/or long-range cellular communication) and it may include potential retransmissions, D_0 is the packet delivery deadline, and P_0 is the minimum probability of delivering each packet before deadline. This model is particularly suitable for event-triggered Decentralized Environmental Notification Messages (DENMs) and real-time wireless-networked control in general [6]. Alternative approaches capturing attributes such as time between delivery of consecutive packets (a.k.a. update delay [7]) and age of information [8] may be particularly suitable for Cooperative Awareness Messages (CAMs), a.k.a. Basic Safety Messages (BSMs) periodically exchanged among vehicles.

The above models can be extended to capture the time-varying nature of V2X networking. For instance, Model (1) may be extended to the following:

$$\text{Probability}\{D(t) \leq D_0(t)\} \geq P_0(t), \quad (2)$$

which captures time-varying dynamics in both application requirements and V2X networking. In practice, it is important to understand the timescales of dynamics in wireless communications and vehicles and capture their impact on V2X modeling. It will also be interesting to explore how to best use such/similar models in vehicle control and sensing. Besides temporal variations and correlations in V2X communication, it will also be important to understand spatial variations and correlations in V2X networking and explore how to apply them in system design and analysis.

Major contributors to V2X communication delay include medium-access control (MAC) delay (including channel contention but also neighbor/link discovery delay for mmWave), propagation delay, and transmission delay. MAC delay is a factor that can be reduced by many mechanisms, and mechanisms should be explored such as infrastructure-assisted scheduling/access-control, distributed Time Division Multiple Access (TDMA) in highly-dynamic settings, transmit-time-aware sampling, and transmit power control (which impacts conflict sets in channel access). In V2X networks, a common primitive is for a vehicle to share its state/operation with close-by vehicles via broadcast. Thus, effectively capturing and controlling broadcast delay becomes an important issue. In fact, there may exist different notions of broadcast delay depending on applications, and one way of defining broadcast delay is to capture the delay in reaching a certain percentage/subset of receivers (i.e., intended receivers).

For high communication reliability, there also exist a wide range of mechanisms that can be exploited in practice. Examples include the following:

- Using multiple antennas and/or communication systems to leverage the enabled diversity.
- Designing effective physical layer solutions such as error control coding, channel estimation, and so on.
- Designing effective MAC mechanisms such as predictable interference control, infrastructure / network-assisted mode, TDMA, priority-aware scheduling, and power control.
- Designing proactive Automatic Repeat Request (ARQ) mechanisms for broadcast which do not require feedback from (all) receivers and leverages predictable broadcast communication reliability control mechanisms.

For high reliability and low latency in V2X communications, it is also critical to leverage redundancy and diversity provided by heterogeneous wireless media and networks, for instance, microwave, mmWave, visible-light, and free-space optical (FSO) communication, as well as the potential availability of the cellular infrastructure. To effectively leverage heterogeneous V2X networks and transmission media in URLL V2X communication, it is important to consider different application requirements as well as properties (e.g., real-time, throughput, and reliability) of different wireless media/networks. It is also important to control the complexity of the resulting system and to potentially fuse non-coherent information from different communication channels in a holistic manner.

4.1.3 Heterogeneous V2X Networking

As mentioned in Section 4.1.2, different technologies can be leveraged to fulfill the challenging requirements of vehicular applications [4]. The most frequently studied family of wireless access technologies are based on Dedicated Short-Range Communications (DSRC) and IEEE 802.11p. The latter ones are mainly intended to support localized vehicle-to-vehicle (V2V) communications. On the other hand, cellular technologies have been considered to support long-range connectivity with remote entities. Recently, 3GPP has sprinted forward by designing in Release 14 the Cellular V2X (C-V2X) technology, which supports V2V interactions over the PC5 sidelink interface in the 5.9 GHz spectrum. The sidelink interface is expected to further evolve with the New Radio (NR) technology expected to be specified in Release 16 for fifth generation (5G) systems. Outside of cellular and 802.11p-based technologies, many other potential technologies candidate themselves for the support of V2X services [4].

There is a wide consensus on the need of a heterogeneous networking solution combining multiple technologies, while outperforming the behavior of a technology alone [4]. Interesting recent examples can be found in [9] where DSRC messages are used to improve mmWave beamforming procedures, and in [10] where DSRC and VLC are used to boost the performance of platooning applications.

However, whenever multiple networking technologies are combined, the challenge of defining a suitable architecture arises.

4.1.3.1 V2X Applications Taxonomy

Applications conceived for improving passenger safety and comfort are typically classified in categories [2, 3, 4] such as:

- **Safety Applications**, e.g., collision avoidance, vulnerable road users warning;
- **Traffic Efficiency and Management**, e.g., local road traffic information exchange;
- **Infotainment**, e.g., Internet access, on-line streaming services;
- **Remote Diagnostics**, e.g., monitoring of the charging state of an electric vehicle;
- **Cooperative Driving**, e.g., platooning, cooperative maneuvering.

The requirements regarding latency, data rate, and reliability for these classes vary to a large extent – for example, safety-related applications based on a cooperative awareness by local communication often require a latency of less than 100 ms, whereas traffic efficiency or remote diagnostic applications can tolerate latencies in the order of tens of seconds.

While for the first four mentioned application classes products are already on the market, cooperative driving applications are not yet implemented in the field, being uniquely tightened to autonomous vehicle operation. They are a representative example of URLL applications [11, 1], that could benefit from a combination of multiple radio and networking technologies, as already discussed in Section 4.1.2.

4.1.3.2 Challenges in Heterogeneous Networking

Besides fulfilling the performance requirements of the respective applications by blending available networking technologies, several further challenges exist – often caused by the different concepts of the technologies.

Addressing

A common assumption for applications requiring communication in the local area (such as safety applications) is to rely on broadcasting, while long distance communication (e.g., for infotainment or remote diagnostics) is performed via a backend server and based on unicast addresses such as IPv4/IPv6 addresses. However, for cooperative driving some use-cases (e.g., negotiation of trajectories between two vehicles or within a platoon) might require a reliable unicast communication in the local area. In this case, a suitable addressing scheme is required, which may need to be complemented by a proper neighbor discovery approach. Some solutions are currently under discussion within IETF [12], but for 802.11p/WAVE networks, not for the C-V2X technology. Furthermore, the vehicle itself may have different addresses for different wireless technologies, leading to the need for a global identifier such as the Vehicle Identification Number (VIN). However, this global identifier might contradict privacy requirements – which could be solved by using temporary identifies (similar to the Temporary Mobile Subscriber Identity, TMSI, in cellular networks but technology independent).

Message Format


Currently, message formats, e.g. the format of a CAM, are often bound to a specific communication technology, despite their access-neutral design. Within a heterogeneous V2X network, a conversion of message formats as well as an aggregation of several messages coming from different communication interfaces can become necessary.

Network Selection

Each vehicle continuously monitors which communication networks are detected in the local situation. When multiple networks are available, selection criteria such as the network load or Quality of Service (QoS) guarantees need to be applied in order to select the optimal network for the local situation. If a handover to the selected technology is assumed, these criteria are sometimes referred to as handover triggers [13]. Instead of performing a handover, simultaneous usage of multiple technologies or per-packet selection of a communication technology [14] can lead to a better performance. The issue of where the network selection should be enforced has to be addressed (e.g., cloud-assisted hybrid vehicular networking [15] or a completely distributed approach.)

Application Model

Due to the wide range of existing and future V2X applications (as already mentioned in Section 4.1.3.1) it is still an open question which application model(s) should be assumed for V2X networking. Depending on the respective applications, a request-response model, a service-oriented model, or a publish-subscribe approach might be more suitable.



■ **Figure 1** Heterogeneous networking abstraction layer avoiding a redundant implementation of technology selection and message format/address conversion.

4.1.3.3 Vehicular Network Architecture

For vehicular networking, several network architectures have been specified, for example ETSI ITS-G5 in Europe (ETSI EN 302 665), IEEE 1609.0/Wireless Access in Vehicular Environments (WAVE) in the US or ARIB STD-T109 in Japan [13]. Besides these standardization efforts, networking solutions for heterogeneous networks have also attracted a widespread interest in the research community which lead to a large number of publications and research projects (a survey of past and recent research efforts can, e.g., be found in [2]).


When considering hybrid vehicular networks where a vehicle has protocol stacks for several vehicular communication technologies on-board, three variants were discussed in the working group:

Class A Traditionally, in the on-board network of a vehicle, an application runs on a single, dedicated electronic-control unit (ECU). In this traditional approach an application, such as collision avoidance, transmitting and receiving CAMs would have the required protocol stack implemented on its ECU, for example the ITS-G5 stack. As a consequence, a single technology per vehicular networking application is used.

Class B The single communication per application approach (Class A) does not allow to leverage the benefits of combining different technologies for a single application. For example, the icy-road ahead warning application using cellular communication in case no other vehicles are in direct communication range can compensate a low market-penetration situation. This can be particularly important in the phase of market introduction of a communication system such as ETSI ITS-G5 or WAVE. One approach to overcome this restriction is to give an application direct access to multiple communication technologies and let the application decide which is the most appropriate technology for the current situation. Since this technology selection process might depend on information on the current status of the communication system (Section 4.1.3.2), the application needs access to all relevant parameters and to implement suitable message formats for all used technologies.

Class C In order to avoid a redundant implementation of status monitoring, message formats and address conversion, this class of hybrid networks introduces a Heterogeneous Vehicular Networking Abstraction Layer (HVNAL), as illustrated in Fig. 1.

The basic idea of the HVNAL is to hide the complexity of the heterogeneous network



■ **Figure 2** Example illustrating the realization of a vehicular networking architecture implementing a heterogeneous vehicular network abstraction layer.

and to be able to implement V2X applications (to some extent) independent of the detailed knowledge of underlying technologies. Thanks to the aforementioned layer in the V2X network architecture, on the one hand, future communication technologies could be introduced without requiring modifications in the individual V2X applications. On the other hand, novel applications, currently unknown, could be supported on top of existing technologies. However, as illustrated in the example in Fig. 2, with an increasing number of available V2X technologies, the complexity of the abstraction layer increases. Furthermore, it is still an open question – also known from classic Internet architectures – in which way the application requirements can be specified in a standardized format at the service access point between application and abstraction layer.

This could be one reason why often the currently discussed architectures for hybrid vehicular networks proposing a similar approach focus on single aspects, for example on load and resource sharing between cellular and direct/ad-hoc networks as in the system investigated by Zheng et. al. in [16]. Here, a Hybrid Link Layer (HLL) for load and resource sharing between cellular networks and IEEE 802.11p is introduced. An alternative approach could be an overlay protocol layer such as the Hybrid Overlay Protocol (HOP) layer in [14] which uses a concept of context indicators to select communication technologies and additionally provides services for data forwarding and aggregation.

4.1.4 Conclusions and Future Research

The requirements of future V2X and cooperative applications cannot be fulfilled by a single communication technology. Due to the large variance in V2X applications – including those envisioned for cooperative driving demanding ultra-reliable low-latency communication – multiple technologies will be required leading to heterogeneous vehicular communication networks.

Nonetheless the plenty of literature solutions, currently, there is no clear consensus on which approach should be implemented and if a hybrid architecture needs to be standardized or can be vendor-specific as long as the message formats and protocols for the individual communication technologies are standardized.

Future research is required to investigate promising solutions such as innovative URLLC communication techniques and architectures supporting a heterogeneous vehicular network abstraction layer. It is worth observing that the idea of an abstraction layer is getting popular in the networking domain, with one of the most prominent instantiation being the Software-defined Networking (SDN) paradigm. The envisioned architecture could treasure SDN principles, currently investigated also a key solution for vehicular networks [17], and further advance them.

The design of the heterogeneous V2X networking architecture could also take inspiration by 5G systems. They face similar issues in the view of supporting multiple applications with different demands on top of the same but properly customized networking facilities, as for instance envisioned by network slicing solutions.

Overall, what clearly emerges from the breakout discussions is that the peculiarities of V2X applications and their continuous evolutions, especially in terms of strict delivery requirements, would require further efforts from the research community in the design of future-proof networking solutions.

References

- 1 R. Johri, J. Rao, H. Yua, and H. Zhang, “A multi-scale spatiotemporal perspective of connected and automated vehicles: Applications and wireless networking,” *IEEE Intelligent Transportation Systems*, vol. 8, no. 2, pp. 65–73, 2016.
- 2 E. Ndashimye, S. K. Ray, N. I. Sarkar, and J. A. Gutiérrez, “Vehicle-to-infrastructure communication over multi-tier heterogeneous networks: A survey,” *Computer Networks*, vol. 112, pp. 144–166, 2017. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128616303826>
- 3 C. Campolo, A. Molinaro, A. Iera, and F. Menichella, “5G network slicing for vehicle-to-everything services,” *IEEE Wireless Communications*, vol. 24, no. 6, pp. 38–45, 2017.
- 4 Z. MacHardy, A. Khan, K. Obana, and S. Iwashina, “V2X access technologies: Regulation, research, and remaining challenges,” *IEEE Communications Surveys & Tutorials*, 2018.
- 5 H. Zhang, X. Liu, C. Li, Y. Chen, X. Che, F. Lin, L. Y. Wang, and G. Yin, “Scheduling with predictable link reliability for wireless networked control,” *IEEE Transactions on Wireless Communications*, vol. 16, no. 9, pp. 6135–6150, 2017.
- 6 Y. Chen, H. Zhang, N. Fisher, L. Y. Wang, and G. Yin, “Probabilistic per-packet real-time guarantees for wireless networked sensing and control,” *IEEE Transactions on Industrial Informatics*, vol. 14, no. 5, pp. 2133–2145, 2018.
- 7 B. Kloiber, C. Rico-Garcia, J. Härrä, and T. Strang, “Update delay: A new information-centric metric for a combined communication and application level reliability evaluation of cam based safety applications,” 2012.

- 8 S. Kaul, M. Gruteser, V. Rai, and J. Kenney, “Minimizing age of information in vehicular networks,” in *Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2011 8th Annual IEEE Communications Society Conference on*. IEEE, 2011, pp. 350–358.
- 9 J. Choi, V. Va, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, and R. W. Heath, “Millimeter-wave vehicular communication to support massive automotive sensing,” *IEEE Communications Magazine*, vol. 54, no. 12, pp. 160–167, 2016.
- 10 M. Segata, R. L. Cigno, H.-M. M. Tsai, and F. Dressler, “On platooning control using IEEE 802.11 p in conjunction with visible light communications,” in *Wireless On-demand Network Systems and Services (WONS), 2016 12th Annual Conference on*. IEEE, 2016, pp. 1–4.
- 11 5GPP, ERTICO ITS EUROPE, and European Commission, *5G Automotive Vision*, Oct. 2015. [Online]. Available: <https://5g-ppp.eu/wp-content/uploads/2014/02/5G-PPP-White-Paper-on-Automotive-Vertical-Sectors.pdf>
- 12 “IPWAVE ipv6 neighbor discovery for prefix and service discovery in vehicular networks,” March 2018.
- 13 K. Abboud, H. A. Omar, and W. Zhuang, “Interworking of DSRC and cellular network technologies for V2X communications: A survey,” *IEEE transactions on vehicular technology*, vol. 65, no. 12, pp. 9457–9470, 2016.
- 14 S. Gopinath, L. Wischhof, C. Ponikwar, and H.-J. Hof, “Hybrid Solutions for Data Dissemination in Vehicular Networks,” in *Proc. 8th International Wireless Days Conference*, Toulouse, France, Mar. 2016.
- 15 T. Higuchi and O. Altintas, “Leveraging cloud intelligence for hybrid vehicular communications,” in *Intelligent Transportation Systems (ITSC), 2017 IEEE 20th International Conference on*. IEEE, 2017, pp. 15–20.
- 16 K. Zheng, Q. Zheng, P. Chatzimisios, W. Xiang, and Y. Zhou, “Heterogeneous vehicular networking: A survey on architecture, challenges, and solutions,” *IEEE communications surveys & tutorials*, vol. 17, no. 4, pp. 2377–2396, 2015.
- 17 G. Han, M. Guizani, Y. Bi, T. H. Luan, K. Ota, H. Zhou, W. Guibene, and A. Rayes, “Software-defined vehicular networks: Architecture, algorithms, and applications: Part 1,” *IEEE Communications Magazine*, vol. 55, no. 7, pp. 78–79, 2017.

4.2 Sensing and Data Management

Suman Banerjee (University of Wisconsin – Madison, US), Aruna Balasubramanian (Stony Brook University, US), Sonia Heemstra de Groot (TU Eindhoven, NL), Albert Held (Daimler AG – Ulm, DE), Frank Kargl (Universität Ulm, DE), Renato Lo Cigno (University of Trento, IT), Thomas Strang (German Aerospace Center-DLR, DE), Lars Wolf (TU Braunschweig, DE), and Andrea Zanella (University of Padova, IT)

License © Creative Commons BY 3.0 Unported license

© Suman Banerjee, Aruna Balasubramanian, Sonia Heemstra de Groot, Albert Held, Frank Kargl, Renato Lo Cigno, Thomas Strang, Lars Wolf, and Andrea Zanella

Cooperative vehicular systems depend heavily on the ability of each vehicle to sense its neighborhood and use this information to interact with neighboring vehicles and the infrastructure. We anticipate that each vehicle is equipped with multi-modal sensors, e.g., to gather audio-visual data, for ranging and positioning, for inertial measurements, and to infer presence of various objects in the neighborhood in three dimensions. Common examples include LIDAR, RADAR, cameras, inertial measurement units, and more. Using these sensed

data, the vehicles aim to learn about other vehicles, pedestrians, various obstacles, and road signage. Given these goals, a number of challenges come to the fore, that are discussed next.

Given the raw sensed data can have a very high volume, is it practical to upload all such data to central repository for different applications. While archiving such raw data might be useful in certain cases, for most real-time applications it is likely adequate to process such data streams locally and only upload or share vectorized content. However, such a design raises a new question – can data captured by one vehicle and processed locally be trusted by another? This is particularly critical if a second vehicle makes various actuation decisions based on data it receives from non-local sources. Sensed data is often inherently noisy. Further, depending on the quality of the sensors and the nature of processing applied, additional inaccuracies may be introduced. Hence, in situations where a vehicle wants to take an action based on data sourced from a different vehicle, the former might benefit from access to the raw data – especially if the first vehicle is more willing to trust its only processing capabilities to extract valuable information out of the sensors.

A next critical issue arises in understanding data ownership. This is a particularly complex issue as there are many stakeholders possible in the data that is sensed. The vehicle manufacturer, the vehicle owner, the objects being sensed, all may have different claims to the data being sensed. This will potentially impact who can do what with the data. Related to ownership is data privacy. For example, camera-based or LIDAR-based system provides raw input from which various contexts of a vehicle and its neighborhood can be extracted. Many vehicle-based video streams go through common privacy preserving techniques, such as face blurring. However, in some cases greater obfuscation techniques might be necessary. The nature of privacy preserving techniques might depend on the context and applications being considered over the data.

Further, sharing such sensed data begs the question on incentives. What is the incentive to share data between different vehicles, especially when they belong to different manufacturers or fleet owners? Some natural incentives exist for sharing data between vehicles of the same manufacturer, e.g., to allow such vehicles to perform some functions like platooning better. Furthermore, the software and hardware subsystems in such vehicles from the same manufacturer are managed by a single entity and data trust is more practical in such scenarios. It is also possible that vehicles across manufacturers may share data with each other in scenarios that improve mutual safety, especially if sharing under such scenarios are mandated through regulations. It is also possible to imagine a credit-based architecture that facilitate sharing across vehicles at a broader scale. The role of regulation might also play an important role in this context.

Sharing of data between vehicles and between vehicles and the infrastructure also require appropriate infrastructure support, especially at the edges of the networks. Requirements include suitable processing, storage, and communication channels to facilitate such sharing, especially when latency is critical.

Finally, to facilitate data sensing and sharing, it is important to define appropriate standards that describe the data and perhaps even policies that identify how different entities may utilize such data for different applications. Overall, sensing, sharing, and data management have many unique challenges that require significant further investigation from various technical standpoints.

4.3 New Use Cases

Sinem Coleri Ergen (Koc University – Istanbul, TR), Onur Altintas (TOYOTA InfoTechnology Center USA – Mountain V, US), Ali Balador (RISE SICS – Västerås, SE), Suman Banerjee (University of Wisconsin – Madison, US), Claudia Campolo (University Mediteranea of Reggio Calabria, IT), Falko Dressler (Universität Paderborn, DE), Eylem Ekici (Ohio State University – Columbus, US), Sonia Heemstra de Groot (TU Eindhoven, NL), Geert Heijenk (University of Twente, NL), Renato Lo Cigno (University of Trento, IT), Michele Segata (University of Trento, IT), Christoph Sommer (Universität Paderborn, DE), Jonathan Sprinkle (NSF – Alexandria, US), Andrea Zanella (University of Padova, IT), and Hongwei Zhang (Iowa State University, US)

License © Creative Commons BY 3.0 Unported license

© Sinem Coleri Ergen, Onur Altintas, Ali Balador, Suman Banerjee, Claudia Campolo, Falko Dressler, Eylem Ekici, Sonia Heemstra de Groot, Geert Heijenk, Renato Lo Cigno, Michele Segata, Christoph Sommer, Jonathan Sprinkle, Andrea Zanella, and Hongwei Zhang

The widespread usage of vehicular networking depends highly on developing a large number of use cases. This working group focused on brainstorming new use cases for vehicular networking. We have generated the following list:

1. **Emergency:** Earthquake and other disaster scenarios may destroy the cellular infrastructure. In those cases, cellular communication may not possible and vehicle-to-vehicle communication may be the only option.
2. **City-wide Surveillance:** Vehicular communication can be used to track people. This information can be further used to derive the movement pattern of people.
3. **Detection of bicycles and pedestrians:** A phone application on the bicycle and pedestrians can communicate with the cloud or directly with the vehicles. The vehicles can then collect this information and combine them with sensor data to detect bicycles and pedestrians.
4. **Distributed black box:** Each car can be considered as a black box, which combines sensor and communication data. When there is an event, such as accident, these data can be retrieved from the database to analyze the statistics related to the event.
5. **Socializing:** The drivers within vehicles close to each other can send warning messages or just to say hello to each other.
6. **Enforcing unwritten rules:** In some countries, there are some unwritten rules. For instance, in India, people give right of way to people of higher status. Vehicular communication can be used to enforce these rules.

4.4 Human-in-the-Loop

Falko Dressler (Universität Paderborn, DE), Eylem Ekici (Ohio State University – Columbus, US), Thorsten Hehn (Volkswagen AG – Wolfsburg, DE), Renato Lo Cigno (University of Trento, IT), Christoph Sommer (Universität Paderborn, DE), and Lars Wischhof (Hochschule München, DE)

License  Creative Commons BY 3.0 Unported license

© Falko Dressler, Eylem Ekici, Thorsten Hehn, Renato Lo Cigno, Christoph Sommer, and Lars Wischhof

4.4.1 Human-in-the-Loop

Despite the many advances in the fields of automated driving, vehicular networking, and now cooperative driving, one major component of the next generation transportation systems has often been ignored or only partially considered: the actual user. Human users as part of the technical system have a significant impact on the design and efficacy of such systems. A novel research domain incorporating human interactions has been termed Cyber Physical Social Systems (CPSS) [1, 2]. In addition to immediate issues related to transitioning from the current road usage to a fully automated one, most prominently the question how to deal with legacy systems, other important societal questions also arise. Considering that human users need to be put first, also to increase public acceptance, the (technical) system must be able to deal with these interactions – hopefully with little impact on efficiency and safety.

4.4.2 Human Beings as a Source of Errors

As cars are slowly moving towards full automation, more functionality is being taken over by the computer. The presence of humans in the decision loop, however, is a source of great uncertainty. This is supported by findings that accurate self-assessment of driving capabilities is massively biased – the well-known Dunning-Kruger effect [3] describes such cognitive bias, wherein persons of low ability suffer from illusory superiority when they mistakenly assess their cognitive ability as greater than it actually is.

So, taking humans out of the control loop has been proposed in our discussion group as a straight forward way of resolving this issue, which would also be the case in the steady state, i.e., when humans feel comfortable with the decisions exclusively made by vehicles. It has been observed that the ownership is one of the major factors in not releasing control of the vehicle – people tend to experience discomfort if *their* car is driven by a computer. However, changing trends in ownership (move towards shared mobility – the ‘Robo-Taxi’ concept) may alleviate some of these issues, leading to public acceptance of automatic driving systems.

It has also been argued that, as long as fully automated driving is not realized, inefficiencies will persist. As an easy way of transitioning to automated driving, back-seat driving options (the former ‘driver’ of the car now acting more as a ‘captain’, ordering a computer what to do) could be used as a transition to full autonomous driving. Similarly, an avatar interacting with the passenger (a natural evolution beyond simple indications of decision processes [4]) would help increase acceptance.

4.4.3 Interfacing Humans and Machines

The interface of computer-driven vehicles with humans and human-controlled vehicles is necessary to ensure harmonious coexistence. This is a longer term issue as it does not depend on acceptance by the driver. The challenges are in replacing interactions with pedestrians

and other human drivers. Such interactions include visual interactions (eye contact, gestures) within appropriate contexts (deployment location, societal norms and habits).

More specifically, interactions with pedestrians are very important: Existing examples include brake lights on the front of the car and projecting crossing lanes for pedestrians. In general, signaling and interacting with pedestrians has been identified as an important topic that need great attention and further research, as pedestrians are still a major components of road casualties. Indeed, research and attention should be extended to all vulnerable road users (VRU) like bicycle riders, but also moped and e-bike users, whose transport mean will not be automated in the foreseeable future. So, adding interfaces (e.g., cell phones) may resolve interaction issues in the short term, and also help adapt to various cultures and environments, but more advanced solutions should also be invented.

4.4.4 Automated Decision Making May Cause Harm

The final discussion was on the moral machine [5, 6], i.e., how to make decisions in critical decision junctures. An existing policy by Volvo, accepting all responsibility in case of an accident, may be relinquishing too much control to the OEM, which might try to reduce the cost rather than implementing other policies. Although automated cars would create a simpler pricing opportunity for insurers, this does not address criminal liability problems. From the human acceptance perspective, it may not be very easy to convince the driver that the controller's decision is better than any real-time decision the driver could have taken and executed.

At this point, we could envision that protecting the driver takes priority in control algorithm design. Protection of others, such as pedestrians, is also taken into account, but as secondary considerations. Legal systems are driven by what is acceptable by society – and legal systems will drive the algorithms controlling the vehicle behavior. Another approach would be to emulate human behavior (possibly including randomness) as a policy.

References

- 1 A. Sheth, P. Anantharam, and C. Henson, "Physical-Cyber-Social Computing: An Early 21st Century Approach," *IEEE Intelligent Systems*, vol. 28, no. 1, pp. 1541–1672, Feb. 2013.
- 2 F. Dressler, "Cyber Physical Social Systems: Towards Deeply Integrated Hybridized Systems," in *IEEE International Conference on Computing, Networking and Communications (ICNC 2018)*. Maui, HI: IEEE, Mar. 2018, pp. 420–424.
- 3 J. Kruger and D. Dunning, "Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments," *Journal of Personality and Social Psychology*, vol. 77, no. 6, pp. 1121–1134, Dec. 1999.
- 4 K. Sonoda and T. Wada, "Displaying System Situation Awareness Increases Driver Trust in Automated Driving," *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 3, pp. 185–193, Sep. 2017.
- 5 J.-F. Bonnefon, A. Shariff, and I. Rahwan, "The social dilemma of autonomous vehicles," *Science*, vol. 352, no. 6293, pp. 1573–1576, Jun. 2016.
- 6 A. Shariff, J.-F. Bonnefon, and I. Rahwan, "Psychological roadblocks to the adoption of self-driving vehicles," *Nature Human Behaviour*, vol. 1, no. 10, pp. 694–696, 2017.

4.5 Safety-critical Vehicular Network Applications

Geert Heijenk (University of Twente, NL), Michele Segata (University of Trento, IT), Christoph Sommer (Universität Paderborn, DE), Thomas Strang (German Aerospace Center-DLR, DE), Lars Wolf (TU Braunschweig, DE), and Andrea Zanella (University of Padova, IT)

License © Creative Commons BY 3.0 Unported license

© Geert Heijenk, Michele Segata, Christoph Sommer, Thomas Strang, Lars Wolf, and Andrea Zanella

4.5.1 Introduction

Today's vehicles are built to very high standards of safety. Embedded systems in the cars, applications running on them, and networks connecting them are rigorously checked to ensure high reliability. As a next step, the safety of drivers will be further enhanced by networks between cars, supporting additional applications that afford the cooperation of multiple cars on the road, e.g., for intersection control or for platooning. However, as vehicle drivers rely more and more on these applications, the emerging behavior of cooperation between cars becomes safety-critical itself: Failure of the cooperation system to perform in a proper manner may directly result in death or serious injury to people. The working group on safety-critical vehicular network applications discussed how, and to what extent, vehicular network applications can be made as close as possible to 100% safe in the presence of faults.

For fully automated cooperative driving systems, foreseen in the future, there appears to be a trade-off between providing functionality and efficiency on the one hand and providing safety on the other hand. At the very extreme, providing 100% safety might mean providing no mobility at all. On the other hand, the fully automated cooperative paradigm can potentially enhance the road safety to levels that may not be otherwise reachable. It is a task of the vehicular networking community to provide insights into this trade-off – and to investigate which level of functionality and performance can be provided at which risk of major failure. Another task is to define measures to minimize the effects of system malfunctioning, especially those due to communication failures.

4.5.2 Comparison to other modes of transport

The working group discussed how the trade-off between performance and risk of failure is made for other modes of transport, especially railways and aviation. A key difference was obvious in the discussion: Whereas the railway fall back to 'fail-safe'(zero velocity) whenever safety requires, an aircraft cannot do the same – it has to remain in some 'fail-operational' mode which allows to continue flying (e.g., go-around maneuver in case of blocked runway). There are many further differences with these transport modes, e.g., the stakeholders, and the incentives for these stakeholders differ, their insights might be applicable to cooperative driving. In railway operation, safety and efficiency seem not to be addressed at the same level. The decision to install a certain safety system along a railway is made almost solely to increase safety, without taking efficiency into consideration too much. Nevertheless, railways do not seem to have a significantly lower capacity per track (12 trains with 800 passengers per hour per track¹ than road systems per lane (2400 cars with 4 passengers capacity per hour [1]). In railway operations, the human train driver is severely limited by what the (safety) system allows him/her to do. Also, in aviation, there seem to be many procedures in

¹ https://en.wikipedia.org/wiki/Route_capacity

place assuring the safety of the aviation as a system. As opposed to railway operation, here it is the pilot who is empowered to carry final responsibility of the aircraft. From another perspective, railway safety is mostly under the control of a single central entity, while in the aviation domain, it is provided in a hierarchical manner, where the aircraft's trajectories are planned and authorized well in advance by a central authority, but can be adjusted during the flight (upon permission) to cope with unpredictable situations, and finally the control can be taken by the pilot in case of immediate danger.

One possible reason for such a different approach to safety in railways and aviation may be the different 'trajectory plasticity' of two scenarios: the trains, indeed, are constrained to follow the train tracks, with basically no possibility of deviation from the planned trajectory, while aircraft can potentially move freely in the 3D space. The plasticity of road vehicle trajectory has its own characteristics, since it is limited by the presence of other nearby vehicles and, obviously, by the road/lane bounds, but can be dynamically and continuously changed within these constraints. Therefore, the approach to road safety in cooperative autonomous driving scenarios may also be different from those considered in railway and aviation scenarios to reflect the specific characteristics of trajectory plasticity of road vehicles.

4.5.3 Directions for solutions

Inspired by the analogy of aviation and railways, two directions for solutions were identified by the working group.

- If the operation of cooperative driving is based on (and depending on) situational awareness by means of sensing and communications, the system (that is, all vehicles in it) should have a good knowledge of the quality of the situational awareness – in terms of accuracy, trustworthiness, freshness, completeness, and correlations in these. Only if the quality of the situational awareness is close to 100%, full functionality and/or performance of the cooperative driving can be employed. If not, the system has to reduce its operation to a less functional, less efficient point of operation. As an example, in the case of platooning operation, headways can be increased depending on the quality of the situational awareness. Of course, degradation time does play an important role here. It should also be taken into account that future autonomous and cooperative systems will not consider the human driver as a possible fail-functional fallback option. As autonomous systems will more and more take over control, humans will progressively lose driving experience. Handing over control to humans might thus increase the likelihood of dangerous situations. This is completely the opposite of what happens in aviation, where pilots are required to manually land the aircraft to keep trained and will resort to automatic landing only if the weather conditions are not good enough. In addition, commercial pilots are required to renew their license every few years. This process is clearly not sustainable for road vehicles, as with autonomous and cooperative driving we are aiming to progressively reduce human intervention.
- Inspired by the railway and aviation scenarios, safety can be considered in a hierarchical way, where all vehicles within a group can operate at very small distances, tightly controlled with highly fault-tolerant operation (compare a series of aircraft in landing configuration on a final glidepath of a runway), whereas safety between groups can be ensured by a combination of less strict coordinated control and larger distances/headways (compare safety between trains). In such a scenario, homogeneity of group members will improve safety, but also introduce dilemmas such as how to enforce homogeneity, e.g., by limiting braking capacity of vehicles in a platoon.

References

- 1 National Research Council – Transportation Research Board, *Highway Capacity Manual 2000: HCM 2000 (metric units)*, 4th ed. TRB, 2000.

4.6 Security and Privacy

Frank Kargl (Universität Ulm, DE), Albert Held (Daimler AG – Ulm, DE), Elmar Schoch (BMW AG – München, DE), Christoph Sommer (Universität Paderborn, DE), Thomas Strang (German Aerospace Center-DLR, DE), Isabel Wagner (De Montfort University – Leicester, GB), and Andrea Zanella (University of Padova, IT)

License © Creative Commons BY 3.0 Unported license
 © Frank Kargl, Albert Held, Elmar Schoch, Christoph Sommer, Thomas Strang, Isabel Wagner, and Andrea Zanella

The breakout group on security and privacy tackled three aspects: First, the role of security in systems engineering of automated connected vehicles (section 4.6.1). Second, questions regarding the privacy of cooperative Intelligent Transportation Systems (ITS) and smart cities (section 4.6.2). Third, security perimeters and attack vectors in automated connected driving (section 4.6.3).

4.6.1 The Role of Security in Systems Engineering

The breakout group started its discussions by tackling one of the fundamental questions: *Why do we require a car to be secure?* Discussions quickly arrived at an interesting angle: dependability. In brief, users must be able to trust that the car does what it is supposed to do. Any attack that does not impact dependability is likely to receive little attention from users – not unlike how users are perfectly content with having their personal computers participate in bot nets (attacking servers, sending spam, ...) as long as their own performance is not impacted. Following this reasoning, one might arrive at the realization that (just like home computers) attacks on cooperative automatic cars might be treated by operators as fundamentally unavoidable. All that might be needed is ensuring that a system remains operational in the face of attacks – though possibly with reduced functionality (e.g., using backhaul control loops). What would be needed, though, are guarantees about the maximum impact of an attack on (safety) application performance (and, as a prerequisite, the ability to quantify the effect of attacks and to discriminate between attacks and failures).

One way towards this might be control theory that natively accounts for security through a true fusion of security engineering and control engineering. In a first step, this might take the shape of a new twist on error modeling: the use of error models that are representative of the effect of security attacks. Further on, it will be necessary to find ‘resilience’ boundaries of control systems to malicious information.

Formalizing the security problems of connected automated vehicles, however, is a complex task. Other than in related work like that of Meadows and Pavlovic [1] (where objectives are often straightforward and questions about a compromise of the system are often a simple, binary decision), some attacks on connected automated vehicles may only affect input data stochastically, with the effects adding up and the goal is to keep some specific processes within given bounds. So formalisms like those of Meadows and Pavlovic [1] would have to be extended to also cover such more complex tasks.

On the plus side, however, cooperative automated vehicles offer the opportunity to ‘offload’ phenomenological detection of attacks to surrounding vehicles. As long as a bare minimum of functionality remains active and untainted in the compromised vehicle, surrounding vehicles could be able to collaboratively issue a *failsafe* command to the vehicle’s controller – an idea not unlike that of an *Air Marshal*.

4.6.2 Privacy of ITS and Smart Cities

The second topic discussed by the breakout group was that of privacy in cooperative ITS and smart cities. One of the key problems here is that location privacy is still not well understood and no well-established privacy metrics are commonly applied to compare solutions [2]. This is compounded by user interface issues: ‘*How can one empower users to take good privacy decisions?*’ is a question that has – to date – no commonly accepted solutions.

In connected fully automated driving, however, the situation changes somewhat. Take, for example, a *Robo-Taxi* scenario. As a driver no longer exists and the identity of passengers is no longer intimately tied to that of the car owner, external attacks like overhearing or license plate recognition do not necessarily reveal information about the identity of passengers.

At the same time, however, the attack surface for internal attacks (by the operator) increases. Here, the amount of data generated, processed and stored will increase dramatically. On the plus side, traditional privacy preserving techniques may apply in this scenario.

Concerns can also be raised about the impact of connected automated vehicles on the privacy of others: With computer vision systems in all such cars (and manufacturers and operators likely recording all data in order to limit their liability) massive amounts of data will be collected about other road users, similar to the infamous Google Street View project. If created data could be bound to strong privacy policies, policy enforcement architectures like investigated by the PRECIOUS project [3] may be a viable option in this scenario to prevent data being abused.

4.6.3 Security Perimeters and Attack Vectors

Work in the breakout group concluded with a consideration of novel attack vectors on cooperative automated vehicles and a discussion on new security perimeter concepts (see Figure 3).

Two identified attack vectors that are specific to cooperative automated vehicles are attacks on sensors and sensor integrity, e.g., by feeding fake Lidar echoes to the car [4] or intelligent spoofing of GPS signals) and attacks on map data (some of which might be crowdsourced). Ways around these attack vectors can include communication of error ranges (in the case of sensors) and automated verification with measurements of many sensors (in the case of map data).

Another issue discussed in the working group was one of attack surfaces: In cooperative driving, many applications like CACC envision the interface to driving functions to be exposed to external entities. As a consequence, the physical boundaries of the vehicle can no longer serve as an isolation perimeter. Rather, semantically-aware input validation will be required. This, however, opens up the problem of misbehavior detection filtering out isolated, only locally-observed, but very relevant information such as about an accident or a small patch of black ice. Workarounds currently considered would be the explicit signaling of ‘Here’s some data, and I know it is hard to believe’, i.e., an *accident flag* in messages – although the impact of this remains unclear.





Figure 3 Arrows marked with a red bar are part of the vehicle's security perimeter and may be subject to novel attacks.

As concerns isolation of security domains inside the vehicle, it can be expected that new EE architectures lead to reduced isolation of formerly distinct components and functions. What would be needed here is dynamic security perimeters that might extend beyond individual vehicles. Examples are *trust groups* with, e.g., a cohort of vehicles such as a platoon as a joint security parameter of all vehicles. This could be complemented by a weaker trust boundary between vehicles within the cohort – albeit this runs into the obvious problem of cooperation across OEM borders.

References

- 1 C. Meadows and D. Pavlovic, "Formalizing physical security procedures," in *Security and Trust Management*, A. Jøsang, P. Samarati, and M. Petrocchi, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 193–208.
- 2 I. Wagner and D. Eckhoff, "Technical Privacy Metrics: A Systematic Survey," *ACM Computing Surveys (CSUR)*, vol. 51, no. 3, pp. 57:1–57:38, June 2018.
- 3 F. Kargl, F. Schaub, and S. Dietzel, "Mandatory enforcement of privacy policies using trusted computing principles," in *Intelligent Information Privacy Management Symposium (Privacy 2010)*. Stanford University, USA: AAAI, March 2010. [Online]. Available: <http://vts.uni-ulm.de/doc.asp?id=7278>
- 4 J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, "Remote attacks on automated vehicles sensors: Experiments on camera and lidar," in *Black Hat Europe*, November 2015. [Online]. Available: <https://www.blackhat.com/docs/eu-15/materials/eu-15-Petit-Self-Driving-And-Connected-Cars-Fooling-Sensors-And-Tracking-Drivers-wp1.pdf>



■ **Figure 4** Extension of the validation process towards city-scale scenarios.

4.7 Simulation, Modeling, and Testing

Christoph Sommer (Universität Paderborn, DE), Wai Chen (China Mobile Research Institute – Beijing, CN), Geert Heijenk (University of Twente, NL), Michele Segata (University of Trento, IT), Jonathan Sprinkle (NSF – Alexandria, US), Erik Ström (Chalmers University of Technology – Göteborg, SE), Isabel Wagner (De Montfort University – Leicester, GB), and Hongwei Zhang (Iowa State University, US)

License © Creative Commons BY 3.0 Unported license

© Christoph Sommer, Wai Chen, Geert Heijenk, Michele Segata, Jonathan Sprinkle, Erik Ström, Isabel Wagner, and Hongwei Zhang

4.7.1 The new scale and dimensions of simulating cooperative mobile systems

Traditionally, system development of (communicating) cars starts with an idea about a component that gets implemented and validated in simulation at many different levels of abstraction. If successful, this component can then move on to lab tests of prototypes, lab tests of a complete car, then field operational tests for certification. This implies that, first, the design process stops at component level and, second, that certification and benchmarking is possible at the level of an individual car or a small group of cars. This process is proven for the development of individual systems (like the traditional communicating car).

For cooperative mobile systems, however, certification and benchmarking according to metrics like *fairness* or *safety* will no longer be possible without considering a large number of cooperating cars – up to city scale trials. In addition to requiring experimentation at scale, these trials will need to be perfectly controlled as well; thus, simulation will emerge as the prime means of both validation and certification of such systems.

As a consequence, the research community will need to find a way towards simulating city scale systems of cars with behavior that is identical to the system under study – ideally, provably so.

Moreover, other than the established approach of employing simulation only at the component or car level, simulation and experimentation will need to be employed for validation (and, even more importantly, for cross-validation) at each step of the composition process (from components, to cars, to individual convoys of cars, to smart cities). Figure 4 shows how the classic approach should be extended: The dashed area highlights the new validation domain that is inherently introduced with cooperative mobile systems.

This is particularly challenging for two main reasons: First, the assumption that a composition of (individually validated) systems can simply be considered valid without further testing is dubious at best. Second, in mixed traffic there will need to be a decidedly *human* component modeled in the system – an aspect that also needs further study.

Aside from this new scale of simulating cooperative mobile systems, simulation will also need to explore new dimensions:

In fully automated systems, simulation and benchmarking can no longer fall back on human behavior in a given situation as the gold standard against which to measure system performance. Ultimately, the objective is that these automated systems will outperform humans, in terms of safety and efficiency of the traffic system. For intermediate performance levels, human behavior could be used as a standard to test against.

Another aspect is that, for performance studies and compliance testing of such systems, typical behavior of a system as complex as a complete smart city (as well as individual, rare events) will need to be simulated using novel metrics: In addition to safety and efficiency (for which approaches have been established in the literature), security, privacy, fairness, and resilience are all qualities of a cooperative mobile system for which metrics will need to be defined and tested with.

Finally, as the level of complexity of the functionality provided and hence the scenarios to be tested is ever increasing, huge amounts of data have to be collected from realistic traffic situations to be able to (re)create testing scenarios. Especially, all data from challenging driving situations, including sensed and communicated data, should be made available for simulation and testing purposes.

4.7.2 Towards reproducible simulation studies

A cross-cutting concern of simulation as a tool for research is ensuring reproducible studies, that is, allowing other researchers to both (1) independently verify the validity of conclusions and to (2) build on the findings of others.

In the early days of small-scale simulation (that is, simulation of just a few aspects of isolated components), the simulation model and its underlying assumptions could well be documented within the few pages of text a scientific paper may allow. With today's simulations encompassing vastly complex systems of multiple components all the way up to trained neural networks and the like, however, writing up a text-form description of this system in a way that allows an interested researcher to reproduce the results (let alone in a timely fashion) has become close to impossible.

The research community will thus need to take the next step in sharing data: Where other fields are simply sharing *result* data (if at all), our community must share *input* data. This data takes two different, complementary forms: First, simulation data, i.e., input traces or training sets, probably collected from real-world challenging driving situations. Second, simulation models and tools – either as full source code of the model, the tool, and all necessary libraries or as a (future-proof) ready-to-run simulation. Such data bundles must also document all the assumptions that have gone into their design and that might restrict their validity (lest other researchers, who might not be domain experts in the particular field, misuse the simulation model, the tool, or other input data).

For the latter step, the sharing of simulation models and tools in a perfectly reproducible form, it might be possible to take a page from the playbook of the DevOps community, who have been creating a wealth of tools for fully automated, reproducible software installations (Docker, ansible, ...) as well as documenting assumptions about their design.

Participants

- Onur Altintas
TOYOTA InfoTechnology Center
USA – Mountain V, US
- Ali Balador
RISE SICS – Västerås, SE
- Aruna Balasubramanian
Stony Brook University – US
- Suman Banerjee
University of Wisconsin –
Madison, US
- Claudia Campolo
University Mediterranea of
Reggio Calabria – IT
- Wai Chen
China Mobile Research Institute –
Beijing, CN
- Sinem Coleri Ergen
Koc University – Istanbul, TR
- Falko Dressler
Universität Paderborn – DE
- Eylem Ekici
Ohio State University –
Columbus, US
- Sonia Heemstra de Groot
TU Eindhoven – NL
- Thorsten Hehn
Volkswagen AG – Wolfsburg, DE
- Geert Heijen
University of Twente – NL
- Albert Held
Daimler AG – Ulm, DE
- Frank Kargl
Universität Ulm – DE
- Florian Klingler
Universität Paderborn – DE
- Renato Lo Cigno
University of Trento – IT
- Jörg Ott
TU München – DE
- Elmar Schoch
BMW AG – München, DE
- Michele Segata
University of Trento – IT
- Christoph Sommer
Universität Paderborn – DE
- Jonathan Sprinkle
NSF – Alexandria, US
- Thomas Strang
German Aerospace Center-DLR –
DE
- Erik Ström
Chalmers University of
Technology – Göteborg, SE
- Isabel Wagner
De Montfort University –
Leicester, GB
- Lars Wischhof
Hochschule München – DE
- Lars Wolf
TU Braunschweig – DE
- Andrea Zanella
University of Padova – IT
- Hongwei Zhang
Iowa State University – US