SCAV'18: Report of the 2nd International Workshop on Safe Control of Autonomous Vehicles*

Mario Gleirscher[†] University of York, UK mario.gleirscher@york.ac.uk

Stefan Kugele Technical University of Munich, Germany stefan.kugele@tum.de

Sven Linker University of Liverpool, UK s.linker@liverpool.ac.uk

ABSTRACT This report summarizes the discussions, open issues, take-away messages, and conclusions of the 2nd SCAV workshop. Keywords Autonomy, verification, control, safety, dependability, reliability.

1. ORGANIZATIONAL SUMMARY

E Retrospect of SCAV'17. Our main goal was to frame the not so \checkmark clearly defined problem of safe autonomous vehicle (AV) control. From the 1^{st} SCAV workshop [3], we were able to determine key challenges as summarized in a report in [2].

SCAV'18. This year, we focused on **applied formal verifica-tion of AV controllers** as well as on discussing their fault-tolerance and assurance. We give an overview of the achievements and identify further issues in the verification of AV control.

including occasional shepherding. We were able to accept 6 out of 8 submissions. The program is outlined in Table 1 and documented in [1]. For SCAV'18,¹ we again asked paper authors to act as discussants for initiating the question & answer sessions after each talk. This way, the discussions appeared more lively because authors get to know each others' works more closely.

Khalil Ghorbal from the INRIA group² for Hybrid Modeling and Contract-Based Design for Multi-physics Embedded Systems gave an excellent keynote on the state-of-the-art of invariant checking and generation for non-linear systems, pointing at current challenges. His presentation on Simulating and Verifying Cyber-Physical Systems (CPSs): Current Challenges and Novel Research Directions shed light on steps to take towards industrial-scale CPSs and provided a smooth transition to our morning session.

Acknowledgments. Latest at this point, we would like to thank all the authors and presenters³ for putting great effort in their talks as well as our discussants and the audience for posing the right questions and pointing out important issues. Furthermore, our program committee⁴ deserves sincere gratitude for their great

³See the list of authors in the workshop proceedings [1].

Table 1: Outlined workshop program	
9:15 - 10:00	Keynote by Khalil Ghorbal
10:30 - 12:30	Morning session: Modeling, verification, simulation (3 talks)
14:00 - 15:30	Afternoon session: Failure analysis, tolerance, safety argu- mentation (3 talks)
16:40 - 17:30	Discussion and tool demo

standards and careful work. Finally, we thank the CPSWeek 2018 organizers for a high-quality scientific event in a beautiful place full of history: Porto in Portugal.

2. MAJOR APPROACHES TO AV SAFETY

In the following, we highlight four important classes of approaches to AV safety—as well applicable to any other type of systems together with some notes offered to be taken away from the workshop contributions and our discussions [1].

When speaking of measures taken for safety assurance, some like to use the term "product-based" to refer to the system under development as the object to be delivered to a customer, and "process-based" to refer to the engineering process or life cycle of this system. We will align the following discussion with this nomenclature.

2.1 Guaranteeing Reachability and Invariance

Here, we view AV safety as the problem of showing for a controlled continuous dynamical physical system or a discrete state transition system that, given an initial set,

- 1. a desired final set is always reached, and
- 2. a desired set (an invariant) will never be left,

through appropriate influence of the controller or decision procedure and, usually, under some additional constraints and the influence of disturbances. The properties 1 and 2 describe what is also called **control stability**. Safety measures taken are, in a strong sense, "product-based" and focus the control application.

From our discussions, we like to mention some state-of-the-art practices and challenges to be solved with good generality:

^{*}Preprint submitted to ACM SIGSOFT SEN.

[†]Supported by the Deutsche Forschungsgemeinschaft (DFG) under the Grant no. GL 915/1-1.

¹See https://scav.in.tum.de.

²See https://team.inria.fr/hycomes.

⁴See the list of members and sub-reviewers in [1].

The Scaling of Reachability Algorithms and Increasing Verification Coverage. Research addresses techniques for set overand under-approximation to reduce "flowpipe"⁵ computation effort. Discretization schemes based on hyper-reals or hyper-dense domains help reducing verification to decidable problems. However, the elaborate example of *multi-lane traffic snapshots* showed us, how hard it is to efficiently cover the space of real-world scenarios coining initial sets and to keep the number of necessary reachability checks low.

The Specification of Controller Performance. In distributed control or with control technology relying on network communication, research is investigating the derivation of lowest bounds of acceptable computational performance (e.g. communication requirements) from given control loop characteristics.

The Reduction of Model Uncertainty and Increasing Verification Confidence. We find it crucial to strengthen discussions on

- identifying uncertainties in models,
- identifying validity errors in models (i.e. errors other than syntactic or semantic inconsistencies), and
- fixing such errors and assuring sufficient model accuracy.

Model parameters are often conservatively estimated by domain experts (e.g. parameters in probabilistic models). Because modeling large systems is challenging, relying solely on expert knowledge of such parameters is not a first class approach to reducing model uncertainty and, hence, can drastically decrease the confidence of any verification results. Clearly, on the other hand, the determination of valid, ideally non-conservative, probabilities of hardware/software failures and similar events is still far from being a trivial issue.

The Transfer of Verification Results to Realistic Settings. Open and flexible simulation platforms for autonomous systems experimentation have shown to be a promising way for this. Particularly, SCAV'18 authors have discussed controller code synthesis from Matlab Simulink⁶ to the ROS⁷ platform and simulation in ROS' RViz environment and the simulation engine Gazebo.⁸ It is not unusual to distinguish between three increasingly more realistic stages of testing by simulation:

- model-in-the-loop (MIL),
- software-in-the-loop (SIL), and
- hardware-in-the-loop (HIL).

Because the assumptions made by these "loop simulations" are increasingly more realistic and, typically, increasingly more expensive, it is important for test and verification engineers to identify the earliest of these stages at which critical aspects of an AV control application can be effectively tested or verified. It is wellknown that late defect discovery is strongly positively correlated with high costs of defect removal.

2.2 Fault-Avoidance and Fault-Tolerance

Here, we view AV safety as the problem of showing for a **controller architecture, design, or implementation** that it fulfills a number of, ideally quantitative, constraints referred to as dependability (incl. reliability) and security requirements. The measures taken are typically "product-based" and focus on controller technology. In practice, many of these constraints are either difficult to identify, quantify, or specify; needless to say that their fulfillment is very difficult in its own right.

Although, our discussions to that extent have been less extensive than in SCAV'17 [2, 3], we like to mention a classical issue investigated in a new technical setting in one of the contributions: For mixed-criticality applications using Ethernet technology and software-defined networking (SDN), it was found that **acceptable fail-over performance in network reconfiguration** is achievable. Our discussion gives reason to believe that recent developments in Ethernet technology may be a suitable option for harmonizing currently heterogeneous and, hence, complex invehicle networks.

An AV design determines the "items" that, across the disciplines involved in AV engineering, define the *contexts* of verification tasks. For example, *functional safety* pertains to electrical, electronic, and software parts of an AV. It is well-known that the decomposition of a system into items correlates with the structure of the organization developing this system. Undesired side-effects of this phenomenon lead to gaps in the reasoning how local results from, for example, functional safety improve the overall safety of a control application. The inverse problem arises when deciding about how a constraint characterizing the control application has to be distributed across parts of a controller implementation.

2.3 Application of Normative Frameworks

Here, we view AV safety as the problem of showing that certain **guidelines**, **policies**, **standards**, **or regulations** (presumably accepted by the corresponding stakeholders and regulatory authorities) have been followed in an engineering process and that **compliance** of this process with such a normative framework is the governing part in the assurance argument. Measures taken are, hence, "process-based," focusing the controller engineering process.

Recent normative frameworks have shown deficiencies for some applications that may have a dangerous impact on compliant products as sold in the markets:

- Fail-operational concepts have been neglected in favor of fail-silent assumptions compatible with more traditional human-in-the-loop settings.
- Agile engineering practices have been postponed in favor of rigid, waterfall-oriented processes, such as the V-model.
- Open source components can not yet be safely integrated with proprietary embedded platforms.

However, in their applied forms, standards often lack clarity (i.e. undesirable universality) in some places and overspecification (i.e. undesirable restriction) in other places. This

⁵Approximation of a continuous dynamics given an initial set.

⁶See https://www.mathworks.com.

⁷See http://www.ros.org.

 $^{^8 \}mathrm{See} \ \mathtt{http://gazebosim.org}.$

circumstance deems them less helpful as a quality control mechanism and in liability and accountability cases. Late availability of the frameworks postpones evaluations of how the mentioned deficiencies are mitigated and, more severely, whether the frameworks comply with the state of the science.

While some frameworks are applicable to very specific types of systems (e.g. four-wheel road vehicles), recent versions try to harmonize regulations and transfer them to systems similar from the viewpoint of functional safety, such as motor cycles. For example, with ISO 26262 (version 2), we were unable to identify the relation of unintended behavior and failure. Most likely, these concepts are intertwined. Our discussion indicates that safety assurance against failures and unintended behavior (also summarized as *malfunctioning behavior*) and safety assurance of nominal behavior should be accomplished in a single coherent framework to identify subtle interference. To this end, it was unclear whether the *safety of the intended function* (SOTIF) standard will address any of the class of approaches according to Section 2.1.

2.4 Case-based Argumentation

Here, we view AV safety as the problem of constructing **an indi**vidual **argument** that the deployed controller is sufficiently free of hazards or, similarly, that the controller is acceptably safe with respect to the identified hazards. This individual argument results from **case-based explicit inductive or deductive reasoning** from evidence towards a desirable claim, or vice versa. The term *evidence* refers to measures of any kind taken to substantiate **or refute** such a claim. The adjective "explicit" is often understood as the hierarchical visualization of this argument. This visualization can, for example, be accomplished by using the Goal Structuring Notation (GSN).⁹ Measures taken can be "product-" or "process-based" and focus both the control application and the controller implementation.

We like to mention an important issue pointed out in our discussion: Safety and, more generally, assurance cases, typically, represent judgments of expert committees about the most probable effects of the taken safety measures. Although such judgments are structured by the argument, "judgmental" uncertainty can well be injected through the evidence and reasoning strategies used to construct the argument. "There is no guarantee" or, in other words, case-based arguments also bear the tremendous challenge of delivering an acceptable level of confidence.

3. COMBINING THESE APPROACHES

For sake of brevity, we decided not to include references in this report. It is worth noting, the body of relevant literature on the presented approaches is overwhelmingly large.

Anyway, the knowledgeable reader might find it quite easy to see that these approaches to AV safety are related. However, the character of these relations and their implications are far less obvious. In summary, the mentioned approaches more or less address two major objectives characterizing the field:

Assurance of Correctness. The approaches described in the Sections 2.1, 2.2 and 2.4 provide capabilities to accomplish the technical part of safe AV control, that is, provide sufficient or compelling evidence that an AV controller implementation actually fulfills its specification, not necessarily questioning the acceptability or validity of this specification. Such a specification

can be seen as a model and such a model, in most of the practical cases, only captures part of the reality in its assumptions. This way, we again refer to the well-known fact that "absolute safety" cannot be guaranteed.

Assurance of Societal Agreeableness. Importantly, in addition to the first objective, the approaches described in the Sections 2.3 and 2.4 provide capabilities for the interaction of system vendors, regulatory authorities, and scientific institutions with public society. A result of such an interaction can, for example, be an agreed definition and *specification* of what it means for an AV controller to be "acceptably safe" or, in other words, of what it means for safety risks in AV control to be "as low as reasonably practicable"¹⁰ in a given operational context.

The workshop talks and our discussions highlighted important steps towards effective combinations of the approaches mentioned in the Sections 2.1 to 2.4:

- We have discussed the integration of approximations of system dynamics with fault models of AV controller implementations. While such approaches are very promising, we believe they will only be of practical use if they can clearly show how to deal with complex fault domains, moreover, how they can effectively relate these domains with the approximated physical loop dynamics.
- Regarding simulation platforms used in practice, we point to the challenge of **combining test and simulation with formal verification** and vice versa.
- We crossed the topic of **complexity and real-time performance** crucial for on-line applications of the discussed verification algorithms.

Conclusion. Our empirical insights and feedback from interviews suggest that practitioners in charge of AV control assurance need to be equipped with and educated in stronger methods than they are currently applying to challenge **assurance scalability**. In conclusion, a **sound combination of the aforementioned approaches** can be seen as a grand challenge and a necessity if we like to push AV control, in particular, and autonomous systems control, in general, to a degree of safety desirable and acceptable in public and domestic spaces. **Automated and integrated formal methods** can play a central role in this context.

4. **REFERENCES**

- M. Gleirscher, S. Kugele, and S. Linker, editors. SCAV'18: Proceedings of the 2nd International Workshop on Safe Control of Autonomous Vehicles, volume 269. EPTCS, 2018. http://eptcs.web.cse.unsw.edu.au/content.cgi?SCAV2018.
- [2] M. Gleirscher, S. Kugele, and J. Sprinkle. Safe control of autonomous & connected vehicles (scav'17) – report from the 1st international workshop at cpsweek 2017. ACM SIGSOFT Software Engineering Notes, 42(3), 2017.
- [3] M. Gleirscher, S. Kugele, and J. Sprinkle, editors. SCAV'17: Proceedings of the 1st International Workshop on Safe Control of Connected and Autonomous Vehicles, New York, NY, USA, 2017. ACM. http://dl.acm.org/citation.cfm?id=3055378.

[4] Health and Safety Executive. *Reducing risks, protecting*

people – HSE's decision-making process. HSE Books, 2001.

⁹See http://www.goalstructuringnotation.info.

¹⁰This is referred to as the ALARP principle [4].