

Generating Safe Autonomous Decision-Making in ROS

Yi Yang

Tom Holvoet

imec-DistriNet
Dept. of Computer Science
KU Leuven
B-3001 Leuven, Belgium

yi.yang@kuleuven.be

tom.holvoet@kuleuven.be

The Robot Operating System (ROS) is a widely used framework for building robotic systems. It offers a wide variety of reusable packages and a pattern for new developments. It is up to developers how to combine these elements and integrate them with decision-making for autonomous behavior. The feature of such decision-making that is in general valued the most is safety assurance.

In this research preview, we present a formal approach for generating safe autonomous decision-making in ROS. We first describe how to improve our existing static verification approach to verify multi-goal multi-agent decision-making. After that, we describe how to transition from the improved static verification approach to the proposed runtime verification approach. An initial implementation of this research proposal yields promising results.

1 Introduction

Autonomous systems have the ability to make decisions to achieve predefined goals without human or other guidance or control. As such, the heart of autonomous systems is their autonomous decision-making. Throughout this article, we use an autonomous mobile robot as an example. This autonomous mobile robot is set to have specific goals, some prior knowledge of the world, and safety specifications. Ideally, it should autonomously make decisions that allow the system to achieve its goals while fully adhering to all provided safety specifications. The most important requirements for autonomous decision-making are effective behaviors and safety assurance.

The Robot Operating System (ROS) [21] is a collection of software libraries and tools for developing robot applications. Due to its distinctive features—modularity, extensive functional libraries, and open source, ROS has become the most well-known framework for constructing robotic systems. In ROS, a number of frequently used functions, including navigation, manipulation, and locomotion, have been carefully developed as reusable packages. These functions can be deployed by engineers, thereby reducing the burden of development. Notably, the most well-developed functions are the implementation of robotic actions, yet there is still significant space for advancement in the decision-making of robots.

The safety of autonomous systems mainly refers to safe decision-making in autonomous systems. For decision-making, we are studying agent programming languages [22]. Compared with other agent programming languages, GOAL [13] is particularly well suited for programming autonomous decision-making, because it generates decisions during execution rather than relying on predefined plan specifications. Formal verification is a convincing approach for providing safety assurance of high-confidence systems [23]. Formal verification can be categorized into static verification and runtime verification. Static verification verifies the given properties of all possible runs of a program, while runtime verification monitors one execution of a program and verifies properties at runtime [1].

This is a research preview that outlines a well-defined research idea of generating safe autonomous decision-making for autonomous systems developed using ROS, combining ROS, decision-making, and

formal verification. The research idea is based on our static verification approach for verifying autonomous decision-making. We have two contributions: first, we explain how to improve the existing static verification approach to verify multi-goal multi-agent decision-making; second, we describe how to transition from the static verification approach to the runtime verification approach. The paper is structured as follows. In Section 2, we introduce the related work to our work. In Section 3, we briefly explain how to improve our static verification approach to verify multi-goal multi-agent autonomous decision-making. In Section 4, we introduce the rationale, the initial design of the architecture of a decision-making node, and the analysis of the proposed runtime verification approach for generating safe autonomous decision-making in ROS. In Section 5, we draw conclusions and indicate future work.

2 Related Work

While model checking is popular in verifying agent programming languages [4] [11], model checking is not a prevalent approach in verifying GOAL. GOAL has many features in common with BDI APLs [8], but it focuses on logical derivation, making it ideal for deductive verification. The inventors of GOAL proposed the first verification framework for GOAL in [3]. This paper demonstrates how to conduct verification for GOAL programs. [16] implemented the verification framework proposed in [3] using Isabelle [20]. Isabelle's implementation provides high confidence for the verification framework, but has limitations: first, it can only specify a subset of GOAL programs, with only a single agent, a single goal, propositional features, and no communication; second, it lacks an automated verification process, so users have to be familiar with Isabelle.

As no well-developed ROS package for generating decision-making is widely accepted by most users, decision-making remains a challenging issue in ROS. In general, there are two approaches to generating decision-making in ROS: providing a decision-making node in ROS, and connecting an agent programming language to ROS.

The conventional approach to generating decision-making in ROS is to provide a decision-making node. Both packages, i.e. [6] and [19], support the generation of decisions in ROS. Both packages do not use formal verification to generate decisions. In other words, neither of these two packages can ensure that ROS will generate safe decisions. Moreover, the decision-making generation procedures in both packages are problem-dependent, i.e., each unique problem necessitates a unique implementation of the decision-making generation process, which adds to the development difficulties for users.

For the purpose of programming complex decision-making, agent programming languages have been developed. Many efforts have been made to provide safety assurance for the generated decisions by agent programming languages [3] [4] [11] [14]. By integrating an agent programming language with ROS, the decision-making in ROS can benefit from the existing research findings on agent programming languages.

[5] presents an interface for the Gwendolen agent programming language [10] to communicate with ROS. Moreover, the AJPF model checker [11] is used to verify the properties of the decisions generated by the Gwendolen language. Therefore, the interface makes the generation of verified decisions feasible in ROS. However, no performance comparison has been conducted between the presented approach in [5] and the traditional approach. Rosbridge [7] is the primary tool used to implement the interface, but Rosbridge is particularly designed for facilitating web applications with ROS [7]. The connection of an agent programming language to ROS is not quite in accordance with the design intent of Rosbridge, which may result in certain performance concerns. [18] describes a generic multi-layer architecture based on ROS-JADE integration for autonomous transport vehicles. The cognitive layer of the proposed

multi-layer architecture is essentially a ROS-JADE node, serving as a communicating function between ROS and JADE [2]. This paper focuses on the generic multi-layer architecture instead of the details of the cognitive layer. The validation of real settings and the efficiency issues remain for their future work. Considering the architectural design of the ROS-JADE node, it lacks safety assurance.

As a lightweight verification approach, runtime verification is usually considered ideal to provide safety assurance for complicated systems, such as robotic applications. Many runtime verification frameworks have been proposed, such as ROSRV [15] and ROSMonitoring [12]. ROSRV and ROSMonitoring provide safety assurance for autonomous systems by intercepting those messages that violate safety requirements.

3 Static Verification

Our prior work [24] serves as the foundation for the proposed research idea. GOAL sets itself apart from the majority of agent programming languages due to the automated decision-making process. This distinctive feature was the primary factor in our decision to adopt it as the instrument to program autonomous decision-making. [24] is the initial proposal for verifying autonomous decision-making by imposing single-goal and single-agent restrictions on the GOAL program. [24] presents a justified algorithm transforming a stratified GOAL program for a single agent and a single goal to its equivalent transition system, equivalent in terms of operational semantics. We now briefly describe how to substantially advanced the work presented in [24] by removing these constraints on GOAL programs. The proposed static verification approach is presumably to verify multi-goal multi-agent decision-making in GOAL.

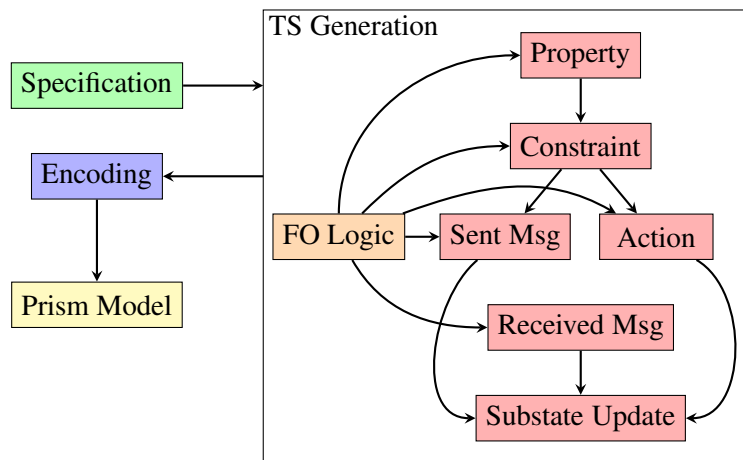


Figure 1: Automated translator for GOAL (Architecture)

Figure 1 shows the architecture of the automated translator. The automated translator consists of two main modules: transition system (TS) Generation and Encoding. The input of the automated translator is the specifications of a multi-agent system that can be specified with GOAL. The output of the automated translator is a Prism model. The automated translator transforms the specifications into its operational-semantically equivalent Prism model that can be directly processed by two efficient probabilistic symbolic model checkers: Storm [9] and Prism [17].

The specification language is an extended version of the specification language provided in [24].

The specification language is inherited from the the formalization of the syntax and semantics of the internal logic of GOAL. The first-order (FO) logic used for the formalization can properly represent the essential characteristics of GOAL programs, including durative actions, multiple goals, multiple agents, and agent communication. The FO logic employed for the formalization serves as a specification language. The specification language can properly specify a multi-agent system in a modular way. We point out that the specification language is suitable for expressing the decision-making generating processes of autonomous systems in addition to GOAL programs. In other words, the specification language can be used to specify the decision-making generating process rather than GOAL. The specifications can be categorized as system-level specifications and agent-level specifications. System-level specifications include knowledge base, communication rules, event processing rules, action generation rules, domains for multi-variables, and constants. Agent-level specifications describe the belief bases and the goal base of an agent.

The crucial step of verifying GOAL programs is the transformation from the GOAL programs to an operational-semantically equivalent transition system. We briefly explain the key concepts of the transformation. In GOAL, the mental state of an agent is a pair consisting of its beliefs and goals [3]. For a GOAL program for a multi-agent system with n agents, each agent has its own mental state. The state of the GOAL program is an n -tuple, whose element is a substate. Precisely, the mental state of an agent is a substate of the GOAL program for a multi-agent system. Each mental state is associated with substate properties. This property is the minimal model of the FO theory consisting of the beliefs and the knowledge base. Based on the logical derivation and the minimal model generation, feasible decisions are generated. A generated decision of a GOAL program is mapped to an action of the transition system. Executed decisions can modify the mental state of an agent. We can conclude that a GOAL program only generates safe decision-making if the generated operational-semantically equivalent transition system satisfies all given safety properties. This is predicated on the assumption that the specification of the GOAL program properly renders the internal logic of the GOAL program. Notably, we use the automated decision-making mechanism to implement the transformation, and the key components of the automated decision-making mechanism are automated logical derivation and minimal model generation. Instead of being problem-dependent, the decision-making generation is logic-dependent, which makes it easier for users to conduct verification.

4 Runtime Verification Approach

4.1 Rationale

The improved static verification approach can provide autonomous systems with the result of PCTL model checking analysis. Assuming that the specifications accurately represent the internal logic of the autonomous decision-making, the static verification approach offers advanced verification results, including convincing safety assurance for the decision-making generation process. However, the static verification approach suffers from certain inherent limitations, such as processing real-time information and the state-space explosion. Based on the existing static verification approach, we propose a runtime verification approach to complement the existing static verification approach. The main component of the proposed runtime verification approach is a decision-making node in ROS. We draw attention to the fact that the runtime verification of autonomous decision-making can benefit from the specification language of the automated translator and the automated decision-making generation.

In summary, Table 1 presents a comparison between the improved static verification approach and the proposed runtime verification approach. Both approaches share the same implementation of the decision-

Comparison	Static Verification	Runtime Verification
DM Generation	Same	Same
Specification language	Same	Same
Specification Provider	User	User and Sensor
State Space	All States	One State
Analysis	PCTL model checking	Safety Checking

Table 1: Comparison

making generation, and they use the same specification language as the input. Users should provide all specifications in the static verification approach, whereas users and sensors provide all specifications in the runtime verification approach. While the runtime verification approach only generates one state at a time, the static verification approach always generates the entire state space. While the static verification approach can perform PCTL model checking, the proposed runtime verification approach only checks the safety at runtime. As the existing static verification approach cannot handle real-time information, this will raise safety concerns at runtime. As a complementary approach, the proposed runtime verification approach focuses on safety checking.

4.2 Architecture for a Decision-Making Node

Figure 2 presents the architecture for the decision-making node in ROS. For a multi-agent system, the decision-making node is expected to generate real-time safe decisions.

Using a multi-agent system consisting of n agents, we demonstrate the similarities and differences between the improved static verification approach and the proposed runtime verification approach.

In terms of the specification language, these system-level specifications are predefined by users in both the static verification approach and the runtime verification approach. However, the agent-level specifications are differently provided in these two approaches. Users and sensors provide the agent-level specifications in the static verification approach and the runtime verification approach, respectively.

The static verification approach and the runtime verification approach share the same implementation of the decision-making generation, but they also contain many modifications to state-space generation and verification. In terms of state-space generation, the entire state space is generated at once via the static verification approach. Specifically, the static method generates every agent decision that could be made to move the system from the initial state to the desired state. The runtime verification approach, on the other hand, only generates one decision at a time, moving from the current state to the following state. In terms of verification, the static verification approach employs probabilistic model checking, using the existing efficient model checkers to verify the properties of the whole state space. The runtime verification approach, on the other hand, merely uses invariant checking while benefiting from intrinsically generated properties during the decision-making generation. As in the static verification approach, each substate is automatically associated with its property. The property checking of the decision-making node can be as simple as checking that all safety specifications are contained in the substate property. The time complexity of this invariant-checking operation is anticipated to be $O(1)$.

Real-time data processing is the key difference between the static verification approach and the runtime verification approach. Real-time data cannot be processed by the static verification method, but they can be processed by the runtime verification approach. Sensors send the decision-making node the real-time beliefs and goals of the agents within the multi-agent system. The decision-making genera-

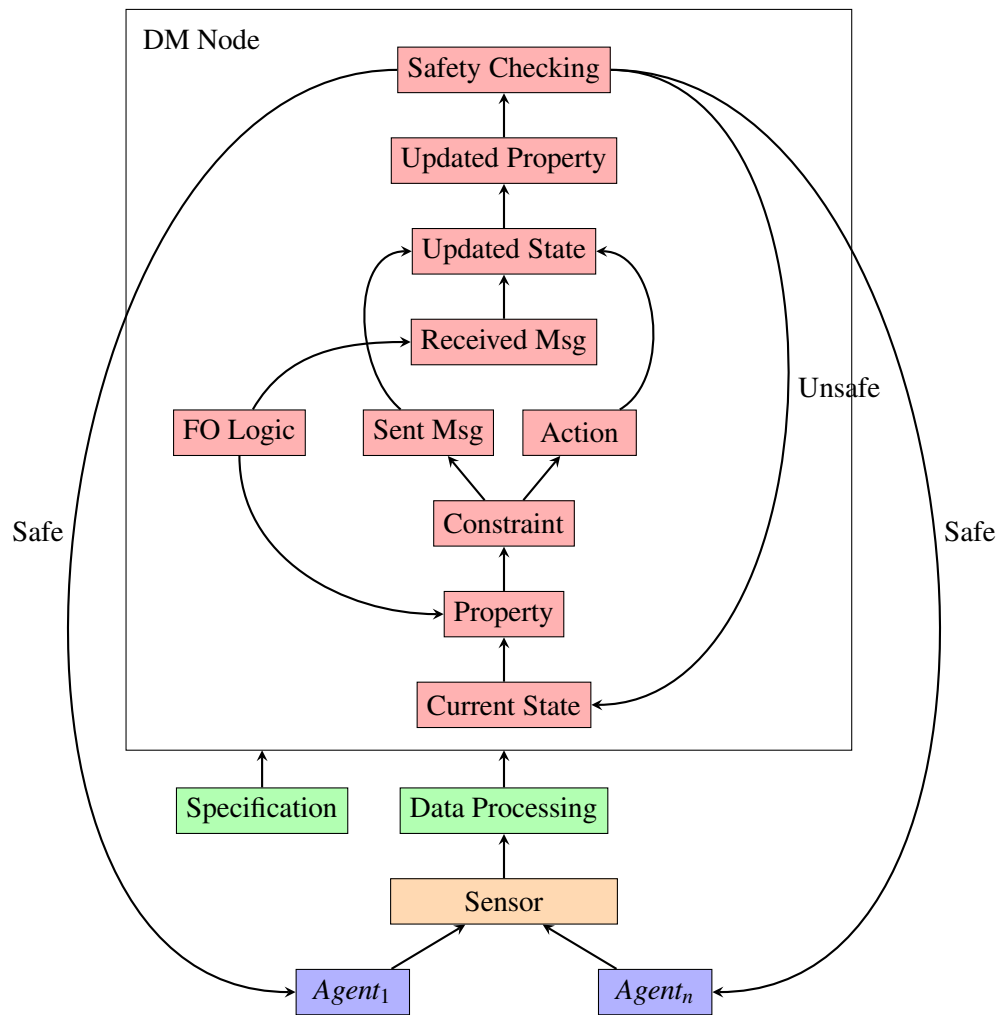


Figure 2: Decision-Making Node (Architecture)

tion module anticipates the ground predicates that the sensor data has been transformed into by the data processing module. The decision-making node generates real-time decisions based on real-time beliefs and goals. Real-time information will presumably be beneficial for decision-making generation for more accuracy.

4.3 Analysis

As the first attempt to transition the static verification approach to the runtime verification approach, the proposed runtime verification approach is expected to be a complementary approach to the improved static verification approach. The main core of the runtime verification approach can hugely benefit from the implementation of the static verification approach. The proposed approach can provide safety assurance at runtime, and it does not suffer from the state-space explosion.

In our setting, each multi-agent system only facilitates one decision-making node, and the decision-making node can only make sequential decisions. Presumably, parallelism could be a bottleneck of the proposed runtime verification approach.

Furthermore, we point out one important issue in the runtime verification approach. We need to generate a transformation between the high-level decision-making and low-level abilities of autonomous agents to enable communication between the decision-making node. The granularity of the decision-making should be properly designed. If the granularity is too high, the decision-making node has to handle too much data, which will incur unnecessary computational costs and reduce the efficiency of generating decisions. If the granularity is too low, the decision-making node has insufficient control over the autonomous agents.

5 Conclusion and Future Work

This paper provides a research preview on how to generate safe autonomous decision-making in ROS. In summary, the proposed runtime verification will be achieved by a two-stage extension of our existing static verification approach. First, we describe how to extend our existing static verification approach to verify multi-goal multi-agent autonomous decision-making. Second, we explain how to transition the improved static verification approach to the proposed runtime verification approach. The main component of the runtime verification approach is a decision-making node. We provided an overview of the architecture of the decision-making node. The decision-making node is expected to generate safe decisions at runtime. We also mentioned that parallelism could be the possible bottleneck of the proposed runtime verification.

We have implemented the initial version of the improved static verification approach, and we validated the initial implementation through two case studies. Our case studies are based on the GOAL source code, provided in the GOAL example project Tower and Coffee of the Eclipse plugin for GOAL. Moreover, we point out that the development of the decision-making node can hugely benefit from the existing implementation of the static verification approach, because they share the same implementation of the decision-making generation. Notably, the static verification approach generates all feasible states from the initial state to the desired state at a time, whereas the runtime verification approach only generates the following state. The static verification approach requires the probabilistic model checking of the whole state space, while the runtime verification approach only requires invariant checking. Therefore, the runtime verification approach is expected to have better performance concerning efficiency than the static verification approach. Initial tests confirm this assumption, yet a thorough analysis will have to validate our hypothesis.

We will fully implement the proposed runtime verification approach in the near future. After that, we will use a multi-agent system consisting of three autonomous mobile robots to validate the improved static verification approach and the runtime verification approach. We also plan to prove the correctness of our implementation of the decision-making node in the future.

Acknowledgements

This research is partially funded by the Research Fund KU Leuven.

References

- [1] Wolfgang Ahrendt, Gordon J Pace & Gerardo Schneider (2012): *A unified approach for static and runtime verification: framework and applications*. In: *International Symposium On Leveraging Applications of Formal Methods, Verification and Validation*, Springer, pp. 312–326, doi:10.1007/978-3-642-34026-0_24.

- [2] Fabio Bellifemine, Federico Bergenti, Giovanni Caire & Agostino Poggi (2005): *JADE—a java agent development framework*. In: *Multi-agent programming*, Springer, pp. 125–147, doi:10.1007/0-387-26350-0_5.
- [3] Frank S de Boer, Koen V Hindriks, Wiebe van der Hoek & J-J Ch Meyer (2007): *A verification framework for agent programming with declarative goals*. *Journal of Applied Logic* 5(2), pp. 277–302, doi:10.1016/j.jal.2005.12.014.
- [4] Rafael H Bordini, Michael Fisher, Carmen Pardavila & Michael Wooldridge (2003): *Model checking agents-peak*. In: *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pp. 409–416, doi:10.1145/860575.860641.
- [5] Rafael C Cardoso, Angelo Ferrando, Louise A Dennis & Michael Fisher (2020): *An interface for programming verifiable autonomous agents in ROS*. In: *Multi-Agent Systems and Agreement Technologies*, Springer, pp. 191–205, doi:10.1007/978-3-642-38700-5_3.
- [6] Cogniteam (2019): *Decision Making Package*. Available at http://wiki.ros.org/decision_making.
- [7] Christopher Crick, Graylin Jay, Sarah Osentoski, Benjamin Pitzer & Odest Chadwicke Jenkins (2017): *Rosbridge: Ros for non-ros users*. In: *Robotics Research*, Springer, pp. 493–504, doi:10.1007/978-3-319-29363-9_28.
- [8] Lavindra De Silva, Felipe Rech Meneguzzi & Brian Logan (2020): *BDI agent architectures: A survey*. In: *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI), 2020, Japão.*, doi:10.24963/ijcai.2020/684.
- [9] Christian Dehnert, Sebastian Junges, Joost-Pieter Katoen & Matthias Volk (2017): *A storm is coming: A modern probabilistic model checker*. In: *International Conference on Computer Aided Verification*, Springer, pp. 592–600, doi:10.1007/978-3-662-49674-9_8.
- [10] Louise A Dennis & Berndt Farwer (2008): *Gwendolen: A BDI language for verifiable agents*. In: *Proceedings of the AISB 2008 Symposium on Logic and the Simulation of Interaction and Reasoning, Society for the Study of Artificial Intelligence and Simulation of Behaviour*, Citeseer, pp. 16–23. Available at <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.141.1549&rep=rep1&type=pdf>.
- [11] Louise A Dennis, Michael Fisher, Matthew P Webster & Rafael H Bordini (2012): *Model checking agent programming languages*. *Automated software engineering* 19(1), pp. 5–63, doi:10.1007/s10515-011-0088-x.
- [12] Angelo Ferrando, Rafael C Cardoso, Michael Fisher, Davide Ancona, Luca Franceschini & Viviana Mascardi (2020): *ROSMonitoring: a runtime verification framework for ROS*. In: *Annual Conference Towards Autonomous Robotic Systems*, Springer, pp. 387–399, doi:10.1023/A:1022920129859.
- [13] Koen V Hindriks (2009): *Programming rational agents in GOAL*. In: *Multi-agent programming*, Springer, pp. 119–157, doi:10.1016/S0004-3702(00)00079-5.
- [14] Gerard J. Holzmann (1997): *The model checker SPIN*. *IEEE Transactions on software engineering* 23(5), pp. 279–295, doi:10.1016/B978-0-444-89874-6.50028-3.
- [15] Jeff Huang, Cansu Erdogan, Yi Zhang, Brandon Moore, Qingzhou Luo, Aravind Sundaresan & Grigore Rosu (2014): *ROSRV: Runtime verification for robots*. In: *International Conference on Runtime Verification*, Springer, pp. 247–254, doi:10.1007/978-3-319-11164-3_20.
- [16] Alexander Birch Jensen, Koen V Hindriks & Jørgen Villadsen (2021): *On Using Theorem Proving for Cognitive Agent-Oriented Programming*. In: *13th International Conference on Agents and Artificial Intelligence*, Science and Technology Publishing, pp. 446–453, doi:10.5220/0010349504460453.
- [17] Marta Kwiatkowska, Gethin Norman & David Parker (2002): *PRISM: Probabilistic symbolic model checker*. In: *International Conference on Modelling Techniques and Tools for Computer Performance Evaluation*, Springer, pp. 200–204, doi:10.1007/3-540-46029-2_13.
- [18] Jon Martin, Oskar Casquero, Brais Fortes & Marga Marcos (2018): *A generic multi-layer architecture based on ROS-JADE integration for autonomous transport vehicles*. *Sensors* 19(1), p. 69, doi:10.1023/A:1020975419546.

- [19] Joao Messias (2015): *Markov Decision Making Package*. Available at http://wiki.ros.org/markov_decision_making.
- [20] Tobias Nipkow, Lawrence C Paulson & Markus Wenzel (2002): *Isabelle/HOL: a proof assistant for higher-order logic*. 2283, Springer Science & Business Media, doi:10.1007/3-540-45949-9_3.
- [21] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, Andrew Y Ng et al. (2009): *ROS: an open-source Robot Operating System*. In: *ICRA workshop on open source software*, 3, Kobe, Japan, p. 5. Available at <http://robotics.stanford.edu/~ang/papers/icraoss09-ROS.pdf>.
- [22] Yoav Shoham (1993): *Agent-oriented programming*. *Artificial intelligence* 60(1), pp. 51–92, doi:10.1016/0004-3702(93)90034-9.
- [23] Jim Woodcock, Peter Gorm Larsen, Juan Bicarregui & John Fitzgerald (2009): *Formal methods: Practice and experience*. *ACM computing surveys (CSUR)* 41(4), pp. 1–36, doi:10.5555/3065491.3065655.
- [24] Yi Yang & Tom Holvoet (2022): *Making Model Checking Feasible for GOAL*. In: *10th International Workshop on Engineering Multi-Agent Systems*. Available at <https://emas.in.tu-clausthal.de/2022/papers/paper4.pdf>.